OS10 Enterprise Edition User Guide

Release 10.4.3.0

DELL EMC

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

© 2019 Dell Inc. or its subsidiaries. All rights reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

2019 - 02

Rev. A00

Contents

1 Getting Started 27
Supported Hardware 27
Download OS10 image and license 28
Installation using ONIE 29
Automatic installation 30
Manual installation 31
Log into OS10 32
Install OS10 license 33
Zero-touch deployment 34
ZTD DHCP server configuration 36
ZTD provisioning script 37
ZTD CLI batch file 37
Post-ZTD script 38
ZTD commands 38
Remote access 40
Configure Management IP address 40
Management Route Configuration 41
Configure user name and password 41
Upgrade OS10 42
CLI Basics 42
User accounts 42
Key CLI features 42
CLI command modes 43
CLI command hierarchy 43
CLI command categories 43
CONFIGURATION Mode 44
Command help 44
Check device status 46
Candidate configuration 48
Change to transaction-based configuration mode 52
Copy running configuration 52
Restore startup configuration 53
Reload system image 53
Filter show commands 54
Alias command 54
Batch mode 58
Linux shell commands 58
SSH commands 59
OS9 environment commands 60
Common commands 60
alias 60

alias (multi-line)....61
batch....62
boot....62
commit....63
configure....63
copy....64
default (alias)....65
delete....65
description (alias)....66
dir....66
discard....67
do....67
feature config-os9-style....68
exit....68
license....69
line (alias)....69
lock....70
management route....70
move....71
no....71
reload....72
show alias....72
show boot....73
show candidate-configuration....74
show environment....76
show inventory....76
show ip management-route....77
show ipv6 management-route....77
show license status....78
show running-configuration....78
show startup-configuration....80
show system....82
show version....84
start....84
system....85
system-cli enable....85
system identifier....85
terminal....86
traceroute....86
unlock....87
write....88

2 Interfaces....89
Ethernet interfaces....89
Unified port groups....89
L2 mode configuration....90

L3 mode configuration ... 91
Fibre Channel interfaces ... 91
Management interface ... 93
VLAN interfaces ... 93
User-configured default VLAN ... 94
VLAN scale profile ... 94
Loopback interfaces ... 95
Port-channel interfaces ... 95
Create port-channel ... 96
Add port member ... 96
Minimum links ... 97
Assign Port Channel IP Address ... 97
Remove or disable port-channel ... 97
Load balance traffic ... 97
Change hash algorithm ... 98
Configure interface ranges ... 98
Switch-port profiles ... 99
S4148-ON Series port profiles ... 100
S4148U-ON port profiles ... 101
Configure breakout mode ... 102
Breakout auto-configuration ... 103
Forward error correction ... 104
Energy-efficient Ethernet ... 105
Enable energy-efficient Ethernet ... 105
Clear EEE counters ... 105
View EEE status/statistics ... 106
EEE commands ... 106
View interface configuration ... 109
Digital optical monitoring ... 112
Enable DOM and DOM traps ... 113
Interface commands ... 114
channel-group ... 114
default vlan-id ... 115
description (Interface) ... 115
duplex ... 116
enable dom ... 116
enable dom traps ... 117
feature auto-breakout ... 117
fec ... 117
interface breakout ... 118
interface ethernet ... 118
interface loopback ... 119
interface mgmt ... 119
interface null ... 119
interface port-channel ... 120

interface range....120
interface vlan....121
link-bundle-utilization....121
mode....121
mode l3....123
mtu....123
port-group....124
scale-profile vlan....124
show interface....124
show inventory media....126
show link-bundle-utilization....126
show port-channel summary....127
show port-group....128
show switch-port-profile....128
show system....129
show vlan....129
shutdown....130
speed (Fibre Channel)....130
speed (Management)....131
switch-port-profile....131
switchport access vlan....133
switchport mode....134
switchport trunk allowed vlan....134

3 Fibre Channel....135
Terminology....136
Virtual fabric....136
Fibre Channel zoning....139
F_Port on Ethernet....140
Pinning FCoE traffic to a specific port of a port-channel....141
Sample FSB configuration on VLT network....143
Sample FC Switch configuration on VLT network....145
Sample FSB configuration on non-VLT network....146
Sample FC Switch configuration on non-VLT network....148
Multi-hop FIP-snooping bridge....149
Configuration notes....150
Configure multi-hop FSB....150
Verify multi-hop FSB configuration....154
Sample Multi-hop FSB configuration....155
Configuration guidelines....165
F_Port commands....165
fc alias....165
fc zone....165
fc zoneset....166
feature fc....166
member (alias)....166

member (zone)..... 167
member (zoneset)..... 167
show fc alias..... 168
show fc interface-area-id mapping..... 168
show fc ns switch..... 168
show fc zone..... 169
show fc zoneset..... 170
zone default-zone permit..... 171
zoneset activate..... 171
NPG commands..... 172
fc port-mode F..... 172
feature fc npg..... 172
show npg devices..... 172
F_Port and NPG commands..... 173
clear fc statistics..... 173
fcoe 174
name..... 174
show fc statistics..... 175
show fc switch..... 175
show running-config vfabric..... 176
show vfabric..... 176
vfabric..... 177
vfabric (interface)..... 177
vlan..... 177
FIP-snooping commands..... 178
feature fip-snooping..... 178
fip-snooping enable..... 178
fip-snooping fc-map..... 179
fip-snooping port-mode..... 179
FCoE commands..... 180
clear fcoe database..... 180
clear fcoe statistics..... 180
fcoe-pinned-port 181
fcoe max-sessions-per-enodemac..... 181
fcoe priority-bits..... 181
lldp tlv-select dcbxp-appln fcoe..... 182
show fcoe enode..... 182
show fcoe fcf..... 183
show fcoe pinned-port..... 183
show fcoe sessions..... 184
show fcoe statistics..... 184
show fcoe system..... 185
show fcoe vlan..... 185

4 Layer 2..... 186
802.1X..... 186

Port authentication....187
EAP over RADIUS....188
Configure 802.1X....188
Enable 802.1X....189
Identity retransmissions....190
Failure quiet period....191
Port control mode....191
Reauthenticate port....192
Configure timeouts....193
802.1X commands....194
Far-end failure detection....198
Enable FEFD globally....200
Enable FEFD on interface....201
Reset FEFD err-disabled interface....201
Display FEFD information....201
FEFD Commands....202
Link Aggregation Control Protocol....205
Modes....205
Configuration....205
Interfaces....206
Rates....207
Sample configuration....207
LACP fallback....211
LACP commands....213
Link Layer Discovery Protocol....220
Protocol data units....221
Optional TLVs....222
Organizationally-specific TLVs....222
Media endpoint discovery....225
Network connectivity device....225
LLDP-MED capabilities TLV....226
Network policies TLVs....227
Define network policies....227
Packet timer values....228
Disable and re-enable LLDP....228
Disable and re-enable LLDP on management ports....229
Advertise TLVs....230
Network policy advertisement....230
Fast start repeat count....231
View LLDP configuration....231
Adjacent agent advertisements....232
Time to live....233
LLDP commands....234
Media Access Control....245
Static MAC Address....246

MAC Address Table ... 246
Clear MAC Address Table ... 246
MAC Commands ... 247
Multiple Spanning-Tree ... 249
Configure MSTP ... 250
Create instances ... 251
Root selection ... 252
Non-Dell EMC hardware ... 253
Region name or revision ... 253
Modify parameters ... 253
Interface parameters ... 254
EdgePort Forward traffic ... 255
Spanning-tree extensions ... 255
Recover BPDU guard error disabled ports ... 257
Setting spanning-tree link type for rapid state transitions ... 258
MST commands ... 258
Rapid per-VLAN spanning-tree plus ... 270
Load balance and root selection ... 271
Enable RPVST+ ... 271
Select root bridge ... 272
Root assignment ... 273
Loop guard ... 274
Global parameters ... 274
Setting spanning-tree link type for rapid state transitions ... 275
RPVST+ commands ... 275
Rapid Spanning-Tree Protocol ... 283
Enable globally ... 283
Global parameters ... 285
Interface parameters ... 286
Root bridge selection ... 287
EdgePort forward traffic ... 287
Spanning-tree extensions ... 288
Setting spanning-tree link type for rapid state transitions ... 289
RSTP commands ... 290
Virtual LANs ... 296
Default VLAN ... 296
Create or remove VLANs ... 297
Access mode ... 298
Trunk mode ... 299
Assign IP address ... 299
View VLAN configuration ... 300
VLAN commands ... 302
Port monitoring ... 303
Local port monitoring ... 303
Remote port monitoring ... 304

Encapsulated remote port monitoring....306
Flow-based monitoring....307
Remote port monitoring on VLT....308
Port monitoring commands....310

5 Layer 3....315
Virtual routing and forwarding....315
Configure management VRF....315
Configure non-default VRF instances....317
VRF configuration....320
View VRF instance information....324
Static route leaking....325
VRF commands....328
Bidirectional Forwarding Detection....336
BFD session states....337
BFD three-way handshake....338
BFD configuration....338
Configure BFD globally....339
BFD for BGP....339
BFD for OSPF....343
BFD for Static route....348
BFD commands....351
Border Gateway Protocol....357
Sessions and peers....358
Route reflectors....359
Multiprotocol BGP....360
Attributes....360
Selection criteria....360
Weight and local preference....361
Multiexit discriminators....362
Origin....362
AS path and next-hop....363
Best path selection....363
More path support....364
Advertise cost....364
4-Byte AS numbers....365
AS number migration....365
Configure Border Gateway Protocol....366
Enable BGP....366
Configure Dual Stack....369
Configure administrative distance....369
Peer templates....370
Neighbor fall-over....373
Configure password....375
Fast external fallover....376
Passive peering....378

Local AS....378
AS number limit....379
Redistribute routes....380
Additional paths....380
MED attributes....381
Local preference attribute....381
Weight attribute....382
Enable multipath....383
Route-map filters....383
Route reflector clusters....383
Aggregate routes....384
Confederations....385
Route dampening....386
Timers....387
Neighbor soft-reconfiguration....387
BGP commands....388
Equal cost multi-path....421
Load balancing....421
Maximum ECMP groups and paths....425
ECMP commands....426
IPv4 routing....430
Assign interface IP address....430
Configure static routing....432
Address Resolution Protocol....432
IPv4 routing commands....433
IPv6 routing....438
Enable or disable IPv6....438
IPv6 addresses....439
Stateless autoconfiguration....440
Neighbor Discovery....441
Duplicate address discovery....442
Static IPv6 routing....443
IPv6 destination unreachable....443
IPv6 hop-by-hop options....444
View IPv6 information....444
IPv6 commands....444
Open shortest path first....457
Autonomous system areas....457
Areas, networks, and neighbors....458
Router types....458
Designated and backup designated routers....459
Link-state advertisements....460
Router priority....461
Shortest path first throttling....461
OSPFv2....462

OSPFv3....495
Object tracking manager....516
Interface tracking....517
Host tracking....517
Set tracking delays....518
Object tracking....518
View tracked objects....519
OTM commands....519
Policy-based routing....522
Policy-based route-maps....522
Access-list to match route-map....522
Set address to match route-map....523
Assign route-map to interface....523
View PBR information....523
Policy-based routing per VRF....524
Configuring PBR per VRF....524
Sample configuration....525
Track route reachability....526
View PBR configuration....526
PBR commands....527
Virtual Router Redundancy Protocol....530
Configuration....530
Create virtual router....531
Group version....532
Virtual IP addresses....532
Configure virtual IP address....533
Configure virtual IP address in a VRF....534
Set group priority....535
Authentication....535
Disable preempt....536
Advertisement interval....536
Interface/object tracking....537
Configure tracking....538
VRRP commands....539

6 Multicast....545
Important notes....545
Configure multicast....545
Unknown multicast flood control....546
Enable multicast flood control....547
Multicast Commands....547
multicast snooping flood-restrict....547
Internet Group Management Protocol....548
Standards compliance....548
Important notes....549
Supported IGMP versions....549

Query interval....549
Last member query interval....549
Maximum response time....549
IGMP immediate leave....550
Select an IGMP version....550
View IGMP-enabled interfaces and groups....550
IGMP snooping....551
IGMP commands....553
Multicast Listener Discovery Protocol....564
MLD snooping....564
MLD snooping commands....566
Protocol Independent Multicast....572
PIM terminology....572
Standards compliance....573
PIM-SM....573
PIM-SSM....574
Expiry timers for S, G entries....574
Static rendezvous point....575
Designated router....575
PIM commands....575
PIM-SM sample configuration....583
PIM-SSM sample configuration....588
Multicast VRF sample configuration....592
Multicast support on VLT....601

7 VXLAN....607
VXLAN concepts....607
VXLAN as NVO solution....608
Configure VXLAN....609
Configure source IP address on VTEP....609
Configure a VXLAN virtual network....610
Configure VLAN-tagged access ports....610
Configure untagged access ports....611
Enable overlay routing between virtual networks....612
Advertise VXLAN source IP address....614
Configure VLT....614
L3 VXLAN route scaling....615
DHCP relay on VTEPs....616
View VXLAN configuration....617
VXLAN MAC addresses....619
VXLAN commands....621
hardware overlay-routing-profile....621
interface virtual-network....622
ip virtual-router address....623
ip virtual-router mac-address....623
member-interface....623

nve....624
remote-vtep....624
show hardware overlay-routing-profile mode....625
show interface virtual-network....625
show nve remote-vtep....626
show nve remote-vtep counters....627
show nve vxlan-vni....627
show virtual-network....627
show virtual-network counters....628
show virtual-network interface counters....628
show virtual-network interface....629
show virtual-network vlan....630
show vlan (virtual network)....630
source-interface loopback....631
virtual-network....631
virtual-network untagged-vlan....632
vxlan-vni....632
VXLAN MAC commands....632
clear mac address-table dynamic nve remote-vtep....632
clear mac address-table dynamic virtual-network....633
show mac address-table count extended....634
show mac address-table count nve....634
show mac address-table count virtual-network....635
show mac address-table extended....635
show mac address-table nve....636
show mac address-table virtual-network....637
Example: VXLAN with static VTEP....638
VTEP 1 Leaf Switch....639
VTEP 2 Leaf Switch....642
VTEP 3 Leaf Switch....645
VTEP 4 Leaf Switch....647
Spine Switch 1....650
Spine Switch 2....651
BGP EVPN for VXLAN....651
BGP EVPN compared to static VXLAN....652
VXLAN BGP EVPN operation....652
Configure BGP EVPN for VXLAN....655
VXLAN BGP EVPN routing....659
BGP EVPN with VLT....659
VXLAN BGP commands....661
VXLAN EVPN commands....664
Example: VXLAN with BGP EVPN....671
Controller-provisioned VXLAN....691
Configuration notes....691
VXLAN components....692

NVO architecture....692
OS10 switch and NSX controller integration....692
Configure the NSX controller....693
Configure and control VXLAN from VMware vCenter....697
Example: VXLAN with a controller configuration....700
VXLAN Controller commands....702

8 UFT modes....709
Configure UFT modes....711
IPv6 extended prefix routes....711
UFT commands....712
hardware forwarding-table mode....712
hardware l3 ipv6-extended-prefix....712
show hardware forwarding-table mode....713
show hardware forwarding-table mode all....713
show hardware l3....713

9 System management....715
Dynamic Host Configuration Protocol....715
Packet format and options....715
DHCP server....717
Automatic address allocation....717
Hostname resolution....718
Manual binding entries....719
DHCP relay agent....720
View DHCP Information....721
System domain name and list....721
DHCP commands....722
DNS commands....728
IPv4 DHCP limitations....730
Network Time Protocol....731
Enable NTP....732
Broadcasts....732
Source IP address....733
Authentication....733
Sample NTP configuration....734
NTP commands....737
System clock....742
System Clock commands....743
System banners....744
Login banner....744
MOTD banner....745
System banner commands....745
User session management....746
User session management commands....747
Telnet server....748

Telnet commands....748
Security....749
User re-authentication....750
Password strength....750
Role-based access control....751
Assign user role....751
Bootloader Protection....752
Linuxadmin User Configuration....752
RADIUS authentication....753
RADIUS over TLS authentication....755
TACACS+ authentication....755
Unknown user role....756
SSH server....757
Virtual terminal line....758
Enable AAA accounting....759
Enable user lockout....759
Limit concurrent login sessions....760
Enable login statistics....760
Privilege levels overview....761
Audit log....763
Security commands....764
X.509v3 certificates....790
Simple Network Management Protocol....808
MIBs....808
SNMP security models and levels....810
SNMPv3....810
SNMP engine ID....810
SNMP groups and users....810
SNMP views....810
Configure SNMP....811
SNMP commands....813
OS10 image upgrade....823
Boot system partition....824
Upgrade commands....824

10 OpenFlow....829
OpenFlow logical switch instance....830
OpenFlow controller....830
OpenFlow version 1.3....830
Ports....830
Flow table....831
Group table....831
Meter table....831
Instructions....831
Action set....832
Action types....832

Counters ... 833
OpenFlow protocol ... 834
OpenFlow use cases ... 848
Configure OpenFlow ... 848
Establish TLS connection ... 849
OpenFlow commands ... 850
controller ... 850
dpid-mac-address ... 851
in-band-mgmt ... 851
max-backoff ... 852
mode openflow-only ... 852
openflow ... 853
probe-interval ... 853
protocol-version ... 853
rate-limit packet_in ... 854
show openflow ... 855
show openflow flows ... 856
show openflow ports ... 856
show openflow switch ... 858
show openflow switch controllers ... 858
switch ... 859
OpenFlow-only mode commands ... 859

11 Access Control Lists ... 862
IP ACLs ... 862
MAC ACLs ... 863
Control-plane ACLs ... 863
Control-plane ACL qualifiers ... 864
IP fragment handling ... 864
IP fragments ACL ... 865
L3 ACL rules ... 865
Permit ACL with L3 information only ... 865
Deny ACL with L3 information only ... 865
Permit all packets from host ... 866
Permit only first fragments and non-fragmented packets from host ... 866
Assign sequence number to filter ... 866
User-provided sequence number ... 866
Auto-generated sequence number ... 866
Delete ACL rule ... 867
L2 and L3 ACLs ... 867
Assign and apply ACL filters ... 868
Ingress ACL filters ... 869
Egress ACL filters ... 869
Clear access-list counters ... 870
IP prefix-lists ... 870
Route-maps ... 871

Match routes.... 872
Set conditions....872
Continue clause....873
ACL flow-based monitoring.... 873
Flow-based mirroring....873
Enable flow-based monitoring.... 874
View ACL table utilization report....875
Known behavior.... 876
ACL logging....877
Important notes....877
ACL commands.... 877
clear ip access-list counters....877
clear ipv6 access-list counters.... 878
clear mac access-list counters.... 878
deny.... 879
deny (IPv6)....879
deny (MAC)....880
deny icmp.... 880
deny icmp (IPv6).... 881
deny ip....881
deny ipv6....882
deny tcp.... 883
deny tcp (IPv6)....883
deny udp.... 884
deny udp (IPv6)....885
description.... 886
ip access-group....886
ip access-list.... 887
ip as-path access-list....887
ip community-list standard deny.... 887
ip community–list standard permit....888
ip extcommunity-list standard deny.... 888
ip extcommunity-list standard permit.... 889
ip prefix-list description....889
ip prefix-list deny....890
ip prefix-list permit....890
ip prefix-list seq deny.... 891
ip prefix-list seq permit.... 891
ipv6 access-group.... 892
ipv6 access-list.... 892
ipv6 prefix-list deny.... 892
ipv6 prefix-list description....893
ipv6 prefix-list permit.... 893
ipv6 prefix-list seq deny.... 894
ipv6 prefix-list seq permit.... 894

mac access-group....895
mac access-list....895
permit....896
permit (IPv6)....896
permit (MAC)....897
permit icmp....897
permit icmp (IPv6)....898
permit ip....898
permit ipv6....899
permit tcp....899
permit tcp (IPv6)....900
permit udp....901
permit udp (IPv6)....902
remark....902
seq deny....903
seq deny (IPv6)....904
seq deny (MAC)....904
seq deny icmp....905
seq deny icmp (IPv6)....906
seq deny ip....907
seq deny ipv6....907
seq deny tcp....908
seq deny tcp (IPv6)....909
seq deny udp....910
seq deny udp (IPv6)....911
seq permit....912
seq permit (IPv6)....913
seq permit (MAC)....913
seq permit icmp....914
seq permit icmp (IPv6)....915
seq permit ip....915
seq permit ipv6....916
seq permit tcp....917
seq permit tcp (IPv6)....918
seq permit udp....919
seq permit udp (IPv6)....920
show access-group....921
show access-lists....922
show acl-table-usage detail....923
show ip as-path-access-list....926
show ip community-list....927
show ip extcommunity-list....927
show ip prefix-list....928
show logging access-list....928
Route-map commands....928

continue ... 929
match as-path ... 929
match community ... 929
match extcommunity ... 930
match interface ... 930
match ip address ... 930
match ip next-hop ... 931
match ipv6 address ... 931
match ipv6 next-hop ... 932
match metric ... 932
match origin ... 932
match route-type ... 933
match tag ... 933
route-map ... 933
set comm-list add ... 934
set comm-list delete ... 934
set community ... 935
set extcomm-list add ... 935
set extcomm-list delete ... 936
set extcommunity ... 936
set local-preference ... 936
set metric ... 937
set metric-type ... 937
set next-hop ... 938
set origin ... 938
set tag ... 939
set weight ... 939
show route-map ... 939

12 Quality of service ... 941
Configure quality of service ... 941
Ingress traffic classification ... 943
Data traffic classification ... 943
Control-plane policing ... 948
Egress traffic classification ... 952
Policing traffic ... 953
Mark Traffic ... 954
Color traffic ... 954
Modify packet fields ... 955
Shaping traffic ... 955
Bandwidth allocation ... 955
Strict priority queuing ... 956
Rate adjustment ... 957
Buffer management ... 958
Configure ingress buffer ... 959
Configure egress buffer ... 960

Deep Buffer mode ... 960
Congestion avoidance ... 962
Storm control ... 963
RoCE for faster access and lossless connectivity ... 963
Configure RoCE on the switch ... 964
Configure RoCE on the interfaces ... 965
RoCE for VXLAN over VLT ... 965
Buffer statistics tracking ... 974
QoS commands ... 975
bandwidth ... 975
buffer-statistics-tracking ... 976
class ... 976
class-map ... 977
clear interface ... 977
clear qos statistics ... 978
clear qos statistics type ... 978
control-plane ... 979
control-plane-buffer-size ... 979
flowcontrol ... 979
hardware deep-buffer-mode ... 980
match ... 980
match cos ... 981
match dscp ... 981
match precedence ... 982
match queue ... 982
match vlan ... 983
mtu ... 983
pause ... 983
pfc-cos ... 984
pfc-max-buffer-size ... 985
pfc-shared-buffer-size ... 985
pfc-shared-headroom-buffer-size ... 985
police ... 986
policy-map ... 986
priority ... 987
priority-flow-control mode ... 987
qos-group dot1p ... 988
qos-group dscp ... 988
qos-rate-adjust ... 988
queue-limit ... 989
queue bandwidth ... 990
queue qos-group ... 990
random-detect (interface) ... 990
random-detect (queue) ... 991
random-detect color ... 991

random-detect ecn....992
random-detect ecn....992
random-detect pool....992
random-detect weight....993
service-policy....993
set cos....994
set dscp....994
set qos-group....994
shape....995
show class-map....995
show control-plane buffers....996
show control-plane buffer-stats....997
show control-plane info....997
show control-plane statistics....998
show hardware deep-buffer-mode....999
show interface priority-flow-control....999
show qos interface....1000
show policy-map....1000
show qos control-plane....1001
show qos egress bufffers interface....1001
show qos egress buffer-statistics-tracking....1002
show qos egress buffer-stats interface....1002
show qos headroom-pool buffer-statistics-tracking....1003
show qos ingress buffers interface....1003
show qos ingress buffer-statistics-tracking....1004
show qos ingress buffer-stats interface....1004
show qos-rate-adjust....1005
show qos service-pool buffer-statistics-tracking....1005
show qos system....1005
show qos system buffers....1006
show qos maps....1007
show qos wred-profile....1009
show queuing statistics....1010
system qos....1010
trust-map....1011
trust dot1p-map....1011
trust dscp-map....1012
qos-map traffic-class....1012
trust-map....1012
wred....1013

13 Virtual Link Trunking....1014
Terminology....1015
VLT domain....1015
VLT interconnect....1016
Graceful LACP with VLT....1016

Configure VLT....1020
MSTP configuration....1021
RSTP configuration....1022
RPVST+ configuration....1023
Create VLT domain....1024
VLTi configuration....1024
Configure VLT MAC address....1025
Delay restore timer....1025
VLT backup....1026
Configure VLT port-channel....1029
VLT unicast routing....1030
VRRP Optimized Forwarding....1030
Migrate VMs across data centers....1031
View VLT information....1035
VLT commands....1037
backup destination....1037
delay-restore....1038
discovery-interface....1038
peer-routing....1038
peer-routing-timeout....1039
primary-priority....1039
show spanning-tree virtual-interface....1040
show vlt....1041
show vlt backup-link....1042
show vlt mac-inconsistency....1042
show vlt mismatch....1043
show vlt role....1045
show vlt vlt-port-detail....1045
vlt-domain....1046
vlt-port-channel....1046
vlt-mac....1046
vrrp mode active-active....1047

14 Uplink Failure Detection....1048
Configure uplink failure detection....1049
Uplink failure detection on VLT....1051
Sample configurations of UFD on VLT....1052
UFD commands....1054
clear ufd-disable....1055
defer-time....1055
downstream....1055
downstream auto-recover....1056
downstream disable links....1056
enable....1056
name....1057
show running-configuration uplink-state-group....1057

show uplink-state-group 1058
uplink-state-group 1059
upstream 1059

15 Converged data center services 1061
Priority flow control 1061
PFC configuration notes 1062
Configure PFC 1064
PFC commands 1067
Enhanced transmission selection 1070
ETS configuration notes 1071
Configure ETS 1071
ETS commands 1073
Data center bridging eXchange 1074
DCBX configuration notes 1074
Configure DCBX 1075
DCBX commands 1077
Internet small computer system interface 1081
iSCSI configuration notes 1082
Configure iSCSI optimization 1082
iSCSI synchronization on VLT 1084
iSCSI commands 1085
Converged network DCB example 1089

16 sFlow 1096
Enable sFlow 1096
Max-header size configuration 1097
Collector configuration 1098
Polling-interval configuration 1098
Sample-rate configuration 1099
Source interface configuration 1100
View sFlow information 1101
sFlow commands 1101
sflow collector 1101
sflow enable 1102
sflow max-header-size 1102
sflow polling-interval 1103
sflow sample-rate 1103
sflow source-interface 1103
show sflow 1104

17 Telemetry 1106
Telemetry terminology 1106
YANG-modeled telemetry data 1106
Configure telemetry 1108
View telemetry configuration 1110

Telemetry commands....1112
debug telemetry....1112
destination....1113
destination-group (subscription-profile)....1113
destination-group (telemetry)....1114
enable....1114
encoding....1114
sensor-group (subscription-profile)....1115
sensor-group (telemetry)....1115
sensor-path....1116
show telemetry....1116
source interface....1118
subscription-profile....1118
telemetry....1118
transport....1119
Example: Configure streaming telemetry....1119

18 RESTCONF API....1121
Configure RESTCONF API....1121
CLI commands for RESTCONF API....1122
rest api restconf....1122
rest https cipher-suite....1122
rest https server-certificate....1123
rest https session timeout....1123
RESTCONF API tasks....1124
Usage Information....1124
View XML structure of CLI commands....1124
RESTCONF API Examples....1125

19 Troubleshoot OS10....1128
Diagnostic tools....1128
Boot partition and image....1129
Monitor processes....1129
LED settings....1130
Packet analysis....1131
Port adapters and modules....1131
Test network connectivity....1132
View solution ID....1133
View diagnostics....1134
Diagnostic commands....1135
Password recovery....1147
Restore factory defaults....1148
SupportAssist....1149
Configure SupportAssist....1149
Set company name....1150
Set contact information....1151

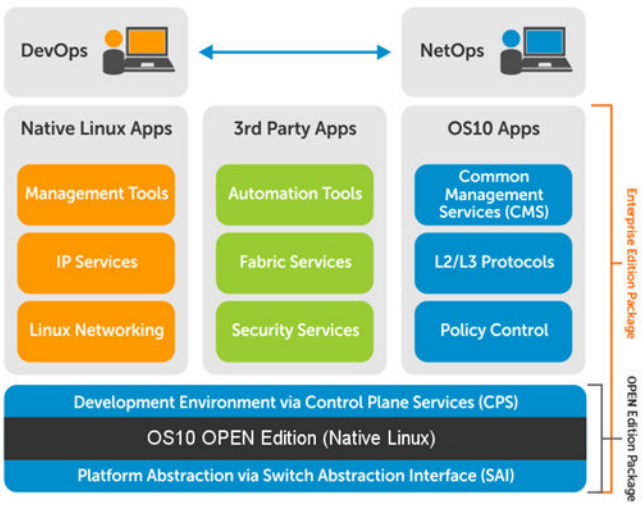
Schedule activity....1151
View status....1152
SupportAssist commands....1153
Support bundle....1159
Event notifications....1160
generate support-bundle....1160
System monitoring....1161
System alarms....1161
System logging....1161
View system logs....1162
Environmental monitoring....1163
Link-bundle monitoring....1163
Alarm commands....1164
Logging commands....1169
Log into OS10 device....1174
Frequently asked questions....1174
Installation....1175
Hardware....1175
Configuration....1175
Security....1176
Layer 2....1176
Layer 3....1176
System management....1176
Access control lists....1177
Quality of service....1177
Monitoring....1178

20 Support resources....1179

1

Getting Started

Dell EMC Networking OS10 Enterprise Edition is a network operating system (OS) supporting multiple architectures and environments. The networking world is moving from a monolithic stack to a pick-your-own-world. The OS10 solution allows disaggregation of the network functionality.



Solutions

- Integrates enabled devices into an existing infrastructure
- Provides up-to-date security fixes that supports a large community of engineers and security experts
- Utilizes an open distribution to simplify customized applications or open source applications

Requirements

- Open network installation environment (ONIE)-enabled Dell EMC device
- OS10 software image stored on an HTTP server or universal serial bus (USB) media
- Familiarity with any Linux release

Supported Hardware

The Dell EMC switches supported are:

- S3048-ON
- S4048-ON, S4048T-ON
- S4112F-ON, S4112T-ON
- S4128F-ON, S4128T-ON
- S4148F-ON, S4148FE-ON, S4148T-ON, S4148U-ON
- S4248FB-ON, S4248FBL-ON
- S5148F-ON

- S5212F-ON, S5224F-ON, S5232F-ON, S5248F-ON, S5296F-ON
- S6010-ON
- Z9100-ON
- Z9264F-ON

NOTE: Starting from release 10.4.2.1, OS10 supports the S5148F-ON platform.

Download OS10 image and license

OS10 Enterprise Edition may come factory-loaded and is available for download from the Dell Digital Locker (DDL). A factory-loaded OS10 image includes a perpetual license. An OS10 image that you download has a 120-day trial license and requires a perpetual license to run beyond the trial period. For more information, see the Setup Guide shipped with your device and the following FAQs:

- Frequently asked Questions
- My Account FAQs

Download an OS10 image and license to:

- Re-install the license on a Dell EMC ONIE switch with a factory-installed OS10 image and license.
- Install OS10 on a Dell EMC ONIE switch without an OS or license installed:
  - A device converted from OS9 or a third-party OS after you uninstall the original OS
  - A replacement device received from Dell EMC return material authorization (RMA)
- Upgrade the OS10 image. For more information, see Upgrade OS10.

Your OS10 purchase allows you to download software images posted within the first 90 days of ownership. To extend the software entitlement, you must have a Dell EMC ProSupport or ProSupport Plus contract on your hardware.

Re-install license on factory-loaded OS10

OS10 Enterprise Edition runs with a perpetual license on an ONIE-enabled device with OS10 factory-loaded. The license file is installed on the switch. If the license becomes corrupted or wiped out, you must download the license from DDL under the purchaser's account and re-install it.

1 Sign in to DDL using your account credentials.
2 Locate the hardware product name with the entitlement ID and order number.
3 Check that the device service tag displays in the `Assigned To:` field on the `Products` page.
4 Click **Key Available for Download**.
5 Select how to receive the license key — by email or downloaded to your local device.
6 Click **Submit**.
7 Save the License.zip file and follow the instructions in Install license to install the license.

Without OS installed

You can purchase the OS10 Enterprise Edition image with an after point-of-sale (APOS) order for a Dell EMC ONIE-enabled device that does not have a default OS or license installed. After the order is complete, you receive an email notification with a software entitlement ID, order number, and link to the DDL.

To extend the entitled download period, bind the software entitlement to the switch service tag to be the same time as the support contract. By default, OS10 software entitlement allows you to download OS10 software images posted before the purchase date and within 90 days of the date.

1 Sign into DDL using your account credentials.
2 Locate your entitlement ID and order number sent by email, then select the product name.
3 On the **Product** page, the `Assigned To:` field on the Product tab is blank. Click **Key Available for Download**.

4 Enter the device service tag you purchased the OS10 Enterprise Edition for in the `Bind to:` and `Re-enter ID:` fields. This step binds the software entitlement to the service tag of the switch.

5 Select how to receive the license key — by email or downloaded to your local device.

6 Click **Submit** to download the License.zip file.

7 Select the **Available Downloads** tab.

8 Select the OS10 Enterprise Edition release to download, then click **Download**.

9 Read the Dell End User License Agreement. Scroll to the end of the agreement, then click **Yes, I agree**.

10 Select how to download the software files, then click **Download Now**.

After you download the OS10 Enterprise Edition image, unzip the .tar file by following these guidelines:

- Extract the OIS10 binary file from the .tar file using any file archiver/compressor software. For example, to unzip a .tar file on a Linux server or from the ONIE prompt, enter:

  ```
  tar -xf tar_filename
  ```

- On a Windows server, some Windows unzip applications insert extra carriage returns (CR) or line feeds (LF) when they extract the contents of a .tar file. The additional CRs or LFs may corrupt the downloaded OS10 binary image. Turn off this option if you use a Windows-based tool to untar an OS10 binary file.
- Generate a checksum for the downloaded OS10 binary image by running the `md5sum` command on the image file. Ensure that the generated checksum matches the checksum extracted from the .tar file.

  ```
  md5sum image_filename
  ```

After you unzip the OS10 Enterprise Edition and download the license, for complete installation and license information, see Installation and Install license.

RMA replacement

A replacement switch comes without an OS or license installed. If you receive a replacement switch, you must assign the STAG of the replacement switch to the SW entitlement in DDL and install the OS10 software and license.

To download OS10 Enterprise Edition and the license, follow the steps for an ONIE switch without an OS installed. For complete installation and license information, see Installation and Install OS10 license.

Installation using ONIE

If you purchase an ONIE-only switch or if you want to replace an existing OS, you can install an OS10 software image using ONIE-based auto-discovery or a manual installation:

- **Automatic installation** — ONIE discovers network information including the Dynamic Host Configuration Protocol (DHCP) server, connects to an image server, and downloads and installs an image automatically.
- **Manual installation** — Manually configure your network information if a DHCP server is not available or if you install the OS10 software image using USB media.

If OS10 is pre-installed on a switch, zero-touch deployment (ZTD) is enabled by default. ZTD automatically downloads and installs an OS10 image in the standby partition. For more information, see Zero-touch deployment.

System setup

Before installation, verify that the system is connected correctly:

- Connect a serial cable and terminal emulator to the console serial port — serial port settings are 115200, 8 data bits, and no parity.
- Connect the Management port to the network to download an image over a network. To locate the Console port and the Management port, see the platform-specific *Installation Guide* at www.dell.com/support.

Install OS10

For an ONIE-enabled switch, navigate to the ONIE boot menu. An ONIE-enabled switch boots up with pre-loaded diagnostics (DIAGs) and ONIE software.

```
+--------------------------------------------------------+
|*ONIE: Install OS                                       |
| ONIE: Rescue                                           |
| ONIE: Uninstall OS                                     |
| ONIE: Update ONIE                                      |
| ONIE: Embed ONIE                                       |
| ONIE: Diag ONIE                                        |
+--------------------------------------------------------+
```

- Install OS — Boots to the ONIE prompt and installs an OS10 image using the Automatic Discovery process. When ONIE installs a new OS image, the previously installed image and OS10 configuration are deleted.
- Rescue — Boots to the ONIE prompt and allows manual installation of an OS10 image or ONIE update.
- Uninstall OS — Deletes the contents of all disk partitions, including the OS10 configuration, except ONIE and diagnostics.
- Update ONIE — Installs a new ONIE version.
- Embed ONIE — Formats an empty disk and installs ONIE.
- EDA DIAG — Runs the system diagnostics.

After the ONIE process installs an OS10 image and you later reboot the switch in `ONIE: Install OS` mode (default), ONIE takes ownership of the system and remains in Install mode (ONIE Install mode is sticky) until an OS10 image successfully installs again. To boot the switch from ONIE for any reason other than installation, select the `ONIE: Rescue` or `ONIE: Update ONIE` option from the ONIE boot menu.

CAUTION: During an automatic or manual OS10 installation, if an error condition occurs that results in an unsuccessful installation and if there is an existing OS on the device, select Uninstall OS to clear the partitions. If the problem persists, contact Dell EMC Technical Support.

Automatic installation

You can automatically install, also known as zero-touch install, an OS10 image on a Dell EMC ONIE-enabled device. After the device successfully boots to ONIE: Install OS, auto-discovery obtains the hostname, domain name, Management interface IP address, and the IP address of the domain name server (DNS) on your network from the DHCP server and DHCP options. The ONIE automatic-discovery process locates the stored software image, starts installation, then reboots the device with the new software image.

If you insert USB device, auto-discovery searches the USB storage supporting FAT or EXT2 file systems. It also searches SCP, FTP, or TFTP servers with the default DNS of the ONIE server. DHCP options are not used to provide the server IP. Auto discovery repeats until a successful software image installation occurs and reboots the switch.

Example for automatic installation

1 Use the `mv image_name onie-installer` command to rename the image as `onie-installer`.

```
mv PKGS_OS10-Base-10.3.1B.144-installer-x86_64.bin onie-installer
```

2 After renaming, the system enters the `ONIE: Install` mode. Enter the command `onie-discovery-start`, which automatically discovers the `onie-installer` image from the DHCP server.

```
ONIE:/ # onie-discovery-start
discover: installer mode detected.  Running installer.
Starting: discover... done.
ONIE:/ # Info: eth0:  Checking link... up.
Info: Trying DHCPv4 on interface: eth0
ONIE: Using DHCPv4 addr: eth0: 10.10.10.17 / 255.0.0.0
Info: eth1:  Checking link... down.
ONIE: eth1: link down.  Skipping configuration.
ONIE: Failed to configure eth1 interface
```

```
ONIE: Starting ONIE Service Discovery
Info: Fetching tftp://10.10.10.2/onie-installer-x86_64-dellemc_s4148fe_c2338 ...
Info: Fetching tftp://10.10.10.2/onie-installer-dellemc_s4148fe_c2338 ...
Info: Fetching tftp://10.10.10.2/onie-installer-x86_64-bcm ...
Info: Fetching tftp://10.10.10.2/onie-installer-x86_64 ...
Info: Fetching tftp://10.10.10.2/onie-installer ...
ONIE: Executing installer: tftp://10.10.10.2/onie-installer
...
...
...
Press <DEL> or <F2> to enter setup.
 Welcome to GRUB!


GNU GRUB  version 2.02~beta2+e4a1fe391
 OS10-B
 EDA-DIAG
 ONIE       Booting `OS10-A'
Loading OS10 ...

[    3.883826] kvm: already loaded the other module
[    3.967628] dummy-irq: no IRQ given.  Use irq=N
[    3.973212] mic_init not running on X100 ret -19
[    3.980168] esas2r: driver will not be loaded because no ATTO esas2r devices were found
[    4.021676] mtdoops: mtd device (mtddev=name/number) must be supplied
[    5.092316] i8042: No controller found
[    5.108356] fmc_write_eeprom fake-design-for-testing-f001: fmc_write_eeprom: no busid
passed, refusing all cards
[    5.120111] intel_rapl: driver does not support CPU family 6 model 77
[    4.226593] systemd-fsck[493]: OS10-SYSROOT1: clean, 23571/426544 files, 312838/1704960
blocks
Debian GNU/Linux 8 OS10 ttyS0
Dell EMC Networking Operating System (OS10)
OS10 login:
```

Manual installation

If a DHCP server is not available, you can manually install an OS10 software image. If the IP address for the Management port (`eth0`) is not automatically discovered, ONIE sets the IP address to `192.168.3.10`. You must manually configure the Management port and configure the software image file to start installation.

1 Save the OS10 software image on an SCP/TFTP/FTP server.

2 Power up the switch and select `ONIE Rescue` for manual installation.

3 (Optional) Stop DHCP discovery if the device boots to ONIE Install.

```
$ onie-discovery-stop
```

4 Configure the IP addresses on the Management port, where *x.x.x.x* represents your internal IP address. After you configure the Management port, the response is `up`.

```
$ ifconfig eth0 x.x.x.x netmask 255.255.0.0 up
```

5 Install the software on the device. The installation command accesses the OS10 software from the specified SCP, TFTP, or FTP URL, creates partitions, verifies installation, and reboots itself.

```
$ onie-nos-install image_filename location
```

For example, enter

```
ONIE:/ # onie-nos-install ftp://a.b.c.d/PKGS_OS10-Enterprise-x.x.xx.bin
```

Where *a.b.c.d* represents the location to download the image file from, and *x.x.xx* represents the version number of the software to install.

The OS10 installer image creates several partitions, including OS10-A (active and default) and OS10-B (standby). After installation completes, the switch automatically reboots and loads OS10.

Install manually using USB drive

You can manually install the OS10 software image using a USB device. Verify that the USB device supports a FAT or EXT2 file system. For instructions to format a USB device in FAT or EXT2 format, see the accompanying Windows documentation for FAT formatting or Linux documentation for FAT or EXT2 formatting.

1 Plug the USB storage device into the USB storage port on the switch.

2 Power up the switch to automatically boot using the `ONIE: Rescue` option.

3 (Optional) Stop ONIE discovery if the device boots to `ONIE: Install`.

```
$ onie-discovery-stop
```

4 Create a USB mount location on the system.

```
$ mkdir /mnt/media
```

5 Identify the path to the USB drive.

```
$ fdisk -l
```

6 Mount the USB media plugged in the USB port on the device.

```
$ mount -t vfat usb-drive-path /mnt/media
```

7 Install the software from the USB, where `/mnt/media` specifies the path where the USB partition is mounted.

```
$ onie-nos-install /mnt/media/image_file
```

The ONIE auto-discovery process discovers the image file at the specified USB path, loads the software image, and reboots the switch. See the ONIE User Guide for more information.

Log into OS10

To log in to OS10 Enterprise Edition, power up the switch and wait for the system to perform a power-on self test (POST). Enter `admin` for both the default user name and user password. For better security, change the default `admin` password during the first OS10 login. The system saves the new password for future logins. After you change the password through the CLI, enter the `write memory` command to save the configuration.

```
OS10 login: admin
Password:
Last login: Sat Oct  6 00:25:33 UTC 2018 on ttyS0
Linux OS10 4.9.110 #1 SMP Debian 4.9.110-3+deb9u4 x86_64

The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.


-*-*-*
```

Install OS10 license

If OS10 is factory-loaded on your switch, you do not need to install an OS10 license. If you download OS10 on a trial basis, OS10 comes with a 120-day trial license. To continue with uninterrupted use, purchase and install a perpetual license to avoid the OS10 device rebooting every 72 hours.

After you install OS10 and log in, install the license to run OS10 Enterprise Edition beyond the trial period. For more information, see Download OS10 image and license. The OS10 license is installed in the */mnt/license* directory.

1 Download the License.zip file from DDL as described in Download OS10 image and license.

2 Open the zip file and locate the license file in the Dell folder. Copy the license file to a local or remote workstation.

3 Install the license file from the workstation in EXEC mode.

```
license install {ftp: | http: | localfs: | scp: | sftp: | tftp: | usb:} filepath/filename
```

- `ftp://userid:passwd@hostip/filepath` — Copy from a remote FTP server.
- `http://hostip/filepath` — Copy from a remote HTTP server.
- `http://hostip` — Send request to a remote HTTP server.
- `localfs://filepath` — Install from a local file directory.
- `scp://userid:passwd@hostip/filepath` — Copy from a remote SCP server.
- `sftp://userid:passwd@hostip/filepath` — Copy from a remote SFTP server.
- `tftp://hostip/filepath` — Copy from a remote TFTP server.
- `usb://filepath` — Install from a file directory on a storage device connected to the USB storage port on the switch.
- *filepath/filename* — Enter the directory path where the license file is stored.

Install license — SCP

```
OS10# license install scp://user:userpwd@10.1.1.10/CFNNX42-NOSEnterprise-License.xml

License installation success.
```

Install license — localfs

Follow these steps to install a license from a local file directory:

1 Copy the license file from the ftp server location to the home directory on the system.

```
OS10# copy ftp://admin:admin@10.11.95.101//home/admin/LADF/7B900Q2-NOSEnterprise-License.XML
home://7B900Q2-NOSEnterprise-License.XML
```

2 (optional) Use the `show copy-file status` command to check the status of the file copy.

```
OS10# show copy-file status
File Transfer State:   idle
--------------------------------------------------
State Detail:          idle
Task Start:            2019-02-15T00:46:35Z
Task End:              2019-02-15T00:46:36Z
Transfer Progress:     100 %
Transfer Bytes:        3795 bytes
File Size:             3795 bytes
Transfer Rate:         8 kbps
```

3 Verify that the license is present in the home directory of your system.

```
OS10# dir home

Directory contents for folder: home
Date (modified) Size (bytes) Name
--------------------- ------------ ----------------------
2019-02-15T00:47:25Z 3795 7B900Q2-NOSEnterprise-License.XML
```

4 Execute the `license install` command specifying the path to the home directory location where the license was downloaded in step 1.

```
OS10# license install localfs://home/admin/7B900Q2-NOSEnterprise-License.XML
[ 5784.994389] EXT4-fs error (device dm-0): ext4_has_uninit_itable:3039: comm
CPS_API_instanc: Inode table for bg 0 marked as needing zeroing
License installation success.
```

Verify license installation

```
OS10# show license status

System Information
---------------------------------------------------------
Vendor Name          :   Dell
Product Name         :   S3048ON
Hardware Version     :   A01
Platform Name        :   x86_64-dell_s3000_c2338-r0
PPID                 :   CN08YWFG282983APSU02
Service Tag          :   7B900Q2
Product Base         :
Product Serial Number:
Product Part Number  :
License Details
----------------
Software        :        OS10-Enterprise
Version         :        10.4.3.0
License Type    :        PERPETUAL
License Duration:        Unlimited
License Status  :        Active
License location:        /mnt/license/7B900Q2.lic
---------------------------------------------------------
```

Troubleshoot license installation failure

An error message displays if the installation fails.

```
License installation failed
```

1 Verify the installation path to the local or remote license location.

2 Check the log on the remote server to find out why the FTP or TFTP file transfer failed.

3 Ping the remote server from the switch — use the `ping` and `traceroute` commands to test network connectivity. Check the following if ping fails:

- If the remote server is reachable through the management route, check if the management route is configured correctly.
- If the remote server is reachable through a front-panel port, check if the static or dynamic route is present.

4 Install the server with the license file on the same subnet as the switch.

5 Check if the server is up and running.

Zero-touch deployment

Zero-touch deployment (ZTD) allows OS10 users to automate switch deployment:

- Upgrade an existing OS10 image.
- Execute a CLI batch file to configure the switch.
- Execute a post-ZTD script to perform additional functions.

ZTD is enabled by default when you boot up a switch with a factory-installed OS10 for the first time or when you perform an `ONIE: OS Install` from the ONIE boot menu. When a switch boots up in ZTD mode, it starts the DHCP client on all interfaces — management and front-panel ports. ZTD configures all interfaces for untagged VLAN traffic. The switch obtains an IP address and a ZTD provisioning script URL from a DHCP server running on the network, and downloads and executes the ZTD script.

- ZTD is supported only in an IPv4 network. ZTD is not supported by DHCPv6.

- At least one of the front-panel ports connected to the network on which the DHCP server is running must be in non-breakout mode.
- After booting up in ZTD mode, if a switch receives no DHCP server response with option 240 within five minutes, it automatically exits ZTD mode. During this time, you can abort ZTD by entering the `ztd cancel` command. The command unlocks the switch configuration so that you can enter OS10 CLI commands.
- When ZTD is enabled, the command-line interface is locked so that you cannot enter OS10 configuration commands. Only show commands are available.

According to the contents of the provisioning script, ZTD performs these tasks in this sequence. Although Steps 2, 3 and 4 are each optional, you must enter a valid URL path for at least one of the `IMG_FILE`, `CLI_CONFIG_FILE`, and `POST_SCRIPT_FILE` variables. For example, if you only want to configure the switch, enter only a `CLI_CONFIG_FILE` URL value. In this case, ZTD does not upgrade the OS10 image and does not execute a post-ZTD script.

1 Downloads the files specified in the ZTD provisioning script — OS10 image, CLI configuration batch file, and post-ZTD script.
   - In the provisioning script, enter the file names for the `IMG_FILE`, `CLI_CONFIG_FILE`, and `POST_SCRIPT_FILE` variables as shown in ZTD provisioning script.
   - If no file names are specified, OS10 immediately exits ZTD and returns to CLI configuration mode.
   - If the download of any of the specified files fails, ZTD stops. OS10 exits ZTD and unlocks the CLI configuration mode.

2 If an OS10 image is specified for `IMG_FILE`, ZTD installs the software image in the standby partition. If no configuration file is specified for `CLI_CONFIG_FILE`, ZTD reloads the switch with the new OS10 image.

3 If an OS10 CLI batch file with configuration commands is specified for `CLI_CONFIG_FILE`, ZTD executes the commands in the `PRE-CONFIG` and `POST-CONFIG` sections. After executing the `PRE-CONFIG` commands, the switch reloads with the new OS10 image and then executes the `POST-CONFIG` commands. For more information, see ZTD CLI batch file.

4 If a post-ZTD script file is specified for `POST_SCRIPT_FILE`, ZTD executes the script. For more information, see Post-ZTD script.

NOTE: The ZTD process performs a single switch reboot. The switch reboot occurs only if either a new OS10 image is installed or if the `PRE-CONFIG` section of the CLI batch file has configuration commands that are executed.

ZTD prerequisites

- Store the ZTD provisioning script on a server that supports HTTP connections.
- Store the OS10 image, CLI batch file, and post-ZTD script on a file server that supports either HTTP, FTP, SFTP, or TFTP connections.
- Configure the DHCP server to provide option 240 that returns the URL of the ZTD provisioning script.
- In the ZTD provisioning script, enter the URL locations of an OS10 image, CLI batch file, and/or post-ZTD script. Enter at least one URL, otherwise the ZTD fails and exits to CLI configuration mode.

ZTD guidelines

- You can store the ZTD provisioning script, OS10 image, CLI batch file, and post-ZTD script on the same server, including the DHCP server.
- Write the ZTD provisioning script in bash.
- Write the post-ZTD script in bash or Python. Enter `#!/bin/bash` or `#!/usr/bin/python` as the first line in the script. The default python interpreter in OS10 is 2.7.

  Use only common Linux commands, such as curl, and common Python language constructs. OS10 only provides a limited set of Linux packages and Python libraries.
- ZTD is disabled by default on automatically provisioned switch fabrics, such as Isilon backend, PowerEdge MX, and VxRail.

Cancel ZTD in progress

To exit ZTD mode and manually configure a switch by entering CLI commands, stop the ZTD process by entering the `ztd cancel` command. You can enter `ztd cancel` only when ZTD is in a waiting state; that is, before it receives an answer from the DHCP server. Otherwise, the command returns an error message; for example:

```
OS10# ztd cancel
% Error: ZTD cancel failed. ZTD process already started and cannot be cancelled at this stage.
```

Disable ZTD

To disable ZTD, enter the `reload` command. The switch reboots in ZTD disabled mode.

Re-enable ZTD

To automatically upgrade OS10 and/or activate new configuration settings, re-enable ZTD by rebooting the switch. Enter the `reload ztd` command. You are prompted to confirm the deletion of the startup configuration.

NOTE: To upgrade OS10 without losing the startup configuration, back up the startup configuration before ZTD runs the provisioning script. Then use the backup startup configuration to restore the previous system configuration.

```
OS10# reload ztd
This action will remove startup-config [confirm yes/no]:
```

View ZTD status

```
OS10# show ztd-status

-----------------------------------
ZTD Status     : disabled
ZTD State      : completed
Protocol State : idle
Reason         : ZTD process completed successfully at Mon Jul 16 19:31:57 2018
-----------------------------------
```

ZTD logs

ZTD generates log messages about its current status.

```
[os10:notify], %Dell EMC (OS10) %ZTD-IN-PROGRESS: Zero Touch Deployment
applying post configurations.
```

ZTD also generates failure messages.

```
[os10:notify], %Dell EMC (OS10) %ZTD-FAILED: Zero Touch Deployment failed to
download the image.
```

Troubleshoot configuration locked

When ZTD is enabled, the CLI configuration is locked. If you enter a CLI command, the error message `configuration is locked` displays. To configure the switch, disable ZTD by entering the `ztd cancel` command.

```
OS10# configure terminal
% Error: ZTD is in progress(configuration is locked).
OS10# ztd cancel
```

ZTD DHCP server configuration

For ZTD operation, configure a DHCP server in the network by adding the required ZTD options; for example:

```
option domain-name "example.org";
option domain-name-servers ns1.example.org, ns2.example.org;
option ztd-provision-url code 240 = text;

default-lease-time 600;
max-lease-time 7200;

subnet 50.0.0.0 netmask 255.255.0.0 {
range 50.0.0.10 50.0.0.254;
option routers rtr-239-0-1.example.org, rtr-239-0-2.example.org;
}

host ztd-leaf1 {
hardware ethernet 90:b1:1c:f4:a9:b1;
fixed-address 50.0.0.8;
option ztd-provision-url "http://50.0.0.1/ztd.sh";
}
```

ZTD provisioning script

Create a ZTD script file that you store on an HTTP server. Configure the URL of the script using DHCP option 240 (`ztd-provision-url`) on the DHCP server.

NOTE: Downloading the ZTD provisioning script is supported only on HTTP connections.

ZTD downloads and executes the script to upgrade the OS10 image, configure the switch, and execute a post-ZTD script to perform additional functions.

- Write the ZTD provisioning script in bash. Enter `#!/bin/bash` as the first line in the script. You can use the sample script in this section as a basis.
- For `IMG_FILE`, enter the URL path of the OS10 image to download and upgrade the switch. The image is written to the standby partition.
- For `CLI_CONFIG_FILE`, enter the URL path of the CLI batch file to download and execute.
- For `POST_SCRIPT_FILE`, enter the URL path of the script to execute.
- ZTD fails and exits to CLI configuration mode if:
  - You do not specify at least one valid URL for the `IMG_FILE`, `CLI_CONFIG_FILE`, and `POST_SCRIPT_FILE` variables.
  - Any of the `IMG_FILE`, `CLI_CONFIG_FILE`, and `POST_SCRIPT_FILE` entries are invalid or if specified, the files cannot be downloaded.

For the `IMG_FILE`, `CLI_CONFIG_FILE`, and `POST_SCRIPT_FILE` files, you can specify HTTP, SCP, SFTP, or TFTP URLs. For example:

```
scp://userid:passwd@hostip/filepath
sftp://userid:passwd@hostip/filepath
```

Example

```
#!/bin/bash

######################################################################
#
#
#             Example OS10 ZTD Provisioning Script
#
#
######################################################################

########## UPDATE THE BELOW CONFIG VARIABLES ACCORDINGLY ###########
########## ATLEAST ONE OF THEM SHOULD BE FILLED ####################

IMG_FILE=”http://50.0.0.1/OS10.bin”
CLI_CONFIG_FILE="http://50.0.0.1/cli_config"
POST_SCRIPT_FILE="http://50.0.0.1/no_post_script.py"


################## DO NOT MODIFY THE LINES BELOW ######################

sudo os10_ztd_start.sh "$IMG_FILE" "$CLI_CONFIG_FILE" "$POST_SCRIPT_FILE"

########################      **END**     ###############################
```

ZTD CLI batch file

Create a CLI batch file that ZTD downloads and executes to configure a switch. The ZTD CLI batch file consists of two sections: `PRE-CONFIG` and `POST-CONFIG`.

ZTD executes the `PRE-CONFIG` commands first using the currently running OS10 image, not the OS10 image specified in the provisioning script. ZTD saves the `PRE-CONFIG` settings to the startup configuration.

If `PRE-CONFIG` commands are present, ZTD reloads the switch before executing the commands in the `POST-CONFIG` section. Enter OS10 configuration commands that require a switch reload, such as `switch-port-profile`, in the `PRE-CONFIG` section. If ZTD installs a new OS10 image (`IMG_FILE`), the new image is activated after the reload.

ZTD then executes the `POST-CONFIG` commands and saves the new settings in the startup configuration. No additional switch reload is performed.

Example

```
# PRE-CONFIG
configure terminal
hostname ZTD-3
exit
configure terminal
interface vlan 210
description ztd-jun29-210
no shutdown
exit

# POST-CONFIG
configure terminal
snmp-server contact DellEMC
exit
configure terminal
interface vlan 500
no shutdown
```

Post-ZTD script

As a general guideline, use a post-ZTD script to perform any additional functions required to configure and operate the switch. In the ZTD provisioning script, specify the post-ZTD script path for the `POST_SCRIPT_FILE` variable. You can use a script to notify an orchestration server that the ZTD configuration is complete. The server can then configure additional settings on the switch.

For example, during the ZTD phase, you can configure only a management VLAN and IP address, then allow an Ansible orchestration server to perform complete switch configuration. Here is a sample curl script that is included in the post-ZTD script to contact an Ansible server:

```
/usr/bin/curl -H "Content-Type:application/json" -k -X POST
--data '{"host_config_key":"'7d07e79ebdc8f7c292e495daac0fe16b'"}'
-u admin:admin https://10.16.134.116/api/v2/job_templates/9/callback/
```

ZTD commands

reload ztd

Reboots the switch and enables ZTD after the reload.

Syntax	`reload ztd`
Parameters	None
Default	ZTD is enabled.
Command Mode	EXEC

Usage Information	Use the `reload ztd` command to automatically upgrade OS10 and/or activate new configuration settings. When you reload ZTD, you are prompted to confirm the deletion of the startup configuration.
Example	`OS10# reload ztd`
Supported Releases	10.4.1.0 or later

show ztd-status

Displays the current ZTD status: enabled, disabled, or canceled.

Syntax	`show ztd-status`
Parameters	None
Default	None
Command Mode	EXEC
Usage Information	None

Examples

```
OS10# show ztd-status

-----------------------------------
ZTD Status     : disabled
ZTD State      : completed
Protocol State : idle
Reason         : ZTD process completed successfully at Mon Jul 16 19:31:57 2018
-----------------------------------
```

```
OS10# show ztd-status

-----------------------------------
ZTD Status     : disabled
ZTD State      : failed
Protocol State : idle
Reason         : ZTD process failed to download post script file
-----------------------------------
```

- `ZTD Status` — Current operational status: enabled or disabled.
- `ZTD State` — Current ZTD state: initialized, in-progress, successfully completed, failed, or canceled while in progress.
- `Protocol State` — Current state of ZTD protocol: initialized, idle while waiting to enable or complete ZTD process, waiting for DHCP post-hook callback, downloading files, installing image, executing pre-config or post-config CLI commands, or executing post-ZTD script file.
- `Reason` — Description of a successful or failed ZTD process.

Supported Releases	10.4.1.0 or later

ztd cancel

Stops ZTD while in progress. After you cancel ZTD, you can enter CLI commands to configure the switch.

Syntax	`ztd cancel`
Parameters	None
Default	ZTD is enabled.
Command Mode	EXEC

Usage Information	When ZTD is enabled, the command-line interface is locked. You cannot enter OS10 configuration commands. Use the `ztd cancel` command to cancel the ZTD process and return to CLI configuration mode. You can enter `ztd cancel` only when ZTD is in a waiting state; that is, before it receives an answer from the DHCP server. Otherwise, the command returns an error message.
Example	`OS10# ztd cancel`
Supported Releases	10.4.1.0 or later

Remote access

You can remotely access the OS10 command-line interface (CLI) and the Linux shell. When you install OS10 the first time, connect to the switch using the serial port.

Configure remote access

- Configure the Management port IP address
- Configure a default route to the Management port
- Configure a user name and password

Remote access OS10 CLI

1 Open an SSH session using the IP address of the device. You can also use PuTTY or a similar tool to access the device remotely.

```
ssh admin@ip-address
password: admin
```

2 Enter `admin` for both the default user name and password to log into OS10. You are automatically placed in EXEC mode.

```
OS10#
```

Remote access Linux shell

```
ssh linuxadmin@ip-address
password: linuxadmin
```

Configure Management IP address

To remotely access OS10, assign an IP address to the management port. The management interface is used for out-of-band (OOB) management purposes.

1 Configure the management interface from CONFIGURATION mode.

```
interface mgmt 1/1/1
```

2 By default, DHCP client is enabled on the Management interface. Disable the DHCP client operations in INTERFACE mode.

```
no ip address dhcp
```

3 Configure an IPv4 or IPv6 address on the Management interface in INTERFACE mode.

```
ip address A.B.C.D/mask
```

```
ipv6 address A:B/prefix-length
```

4 Enable the Management interface in INTERFACE mode.

```
no shutdown
```

Configure Management interface

```
OS10(config)# interface mgmt 1/1/1
OS10(conf-if-ma-1/1/1)# no ip address dhcp
OS10(conf-if-ma-1/1/1)# ip address 10.1.1.10/24
OS10(conf-if-ma-1/1/1)# no shutdown
```

Management Route Configuration

To set up remote access to OS10, configure a management route after you assign an IPv4 or IPv6 address to the Management port. The Management port uses the default management route to communicate with a different network. Management routes are separate from IPv4 and IPv6 routes and are only used to manage the system through the Management port.

```
management route 192.168.100.0/24 1.1.1.1
ip route 192.168.200.0/24 2.2.2.2

management route 192.168.300.0/24 managementethernet
ip route 192.168.400.0/24 interface ethernet 1/1/1
```

Before configuring the static IPv4 address for Management port, remove the DHCP setting using the `no ip address dhcp` command.

Configure a management route to the network in CONFIGURATION mode. Repeat the command to configure multiple routes for the Management port.

```
management route {ipv4-address/mask | ipv6-address/prefix-length}
{forwarding-router-address | managementethernet}
```

- *ipv4-address/mask* — Enter an IPv4 network address in dotted-decimal format (A.B.C.D), then a subnet mask in /prefix-length format (/*x*).
- *ipv6-address/prefix-length* — Enter an IPv6 address in x:x:x:x::x format with the prefix length in /x format. The prefix range is /0 to /128.
- *forwarding-router-address* — Enter the next-hop IPv4/IPv6 address of a forwarding router for network traffic from the Management port.
- `managementethernet` — Configures the Management port as the interface for the route, and associates the route with the Management interface.

Configure management route

```
OS10(config)# management route 10.10.20.0/24 10.1.1.1
OS10(config)# management route 172.16.0.0/16 managementethernet
```

Configure user name and password

To set up remote access to OS10, create a new user name and password after you configure the management port and default route. The user role is a mandatory entry.

Enter the password in clear text. It is converted to SHA-512 format in the running configuration. A password must have at least nine alphanumeric and special characters, and at least five different characters from the password previously used for the same username.

For backward compatibility with OS10 releases 10.3.1E and earlier, passwords entered in MD-5, SHA-256, and SHA-512 format are supported.

To increase the required password strength, use the `password-attributes` command.

- Create a user name and password in CONFIGURATION mode.

  ```
  username username password password role role
  ```

  - `username` *username* — Enter a text string. A maximum of 32 alphanumeric characters; 1 character minimum.
  - `password` *password* — Enter a text string. A maximum of 32 alphanumeric characters; 9 characters minimum.
  - `role` *role* — Enter a user role:
    - `sysadmin` — Full access to all commands in the system, exclusive access to commands that manipulate the file system, and access to the system shell. A system administrator can create user IDs and user roles.

- `secadmin` — Full access to configuration commands that set security policy and system access, such as password strength, AAA authorization, and cryptographic keys. A security administrator can display security information, such as cryptographic keys, login statistics, and log information.
- `netadmin` — Full access to configuration commands that manage traffic flowing through the switch, such as routes, interfaces, and access control lists (ACLs). A network administrator cannot access configuration commands for security features or view security information.
- `netoperator` — Access to EXEC mode to view the current configuration. A network operator cannot modify any configuration setting on a switch.

Create user name and enter password in clear text

```
OS10(config)# username user05 password alpha404! role sysadmin
```

Upgrade OS10

NOTE: During the image upgrade process in a VLT setup, when the VLT peers are running different software versions, no configuration changes should be done in any of the VLT peers. Ensure that both the nodes are upgraded to the same version before you make any configuration change.

To upgrade OS10, download a new OS10 Enterprise Edition image from the DDL.

1 Sign into DDL using your account credentials.
2 Locate the entry for your entitlement ID and order number, then select the product name.
3 Select the **Available Downloads** tab on the Product page.
4 Select the OS10 Enterprise Edition image to download, then click **Download**.
5 Read the Dell End User License Agreement, then scroll to the end of the agreement and click **Yes, I agree**.
6 Select how to download the software files, then click **Download Now**.

Install the OS10 image on an ONIE-enabled switch with an installed OS10 license. For more information, see Install OS10 license.

CLI Basics

The OS10 CLI is the software interface you use to access a device running the software — from the console or through a network connection. The CLI is an OS10-specific command shell that runs on top of a Linux-based OS kernel. By leveraging industry-standard tools and utilities, the CLI provides a powerful set of commands that you can use to monitor and configure devices running OS10.

User accounts

OS10 defines two categories of user accounts — use `admin` for both the username and password to log into the CLI, or use `linuxadmin` to log into the Linux shell.

NOTE: You cannot delete the default admin and linuxadmin usernames.

Key CLI features

Consistent command names	Commands that provide the same type of function have the same name, regardless of the portion of the system on which they are operating. For example, all `show` commands display software information and statistics, and all `clear` commands erase various types of system information.
Available commands	Information about available commands is provided at each level of the CLI command hierarchy. You can enter a question mark (`?`) at any level and view a list of the available commands, along with a short description of each command.

Command completion	Command completion for command names (keywords) and for command options is available at each level of the hierarchy. To complete a command or option that you have partially entered, click the **Tab** key or the **Spacebar**. If the partially entered letters are a string that uniquely identifies a command, the complete command name appears. A beep indicates that you have entered an ambiguous command, and the possible completions display. Completion also applies to other strings, such as interface names and configuration statements.

CLI command modes

The OS10 CLI has two top-level modes:

- EXEC mode — Monitor, troubleshoot, check status, and network connectivity.
- CONFIGURATION mode — Configure network devices.

When you enter CONFIGURATION mode, you are changing the current operating configuration, called the *running configuration*. By default, all configuration changes are automatically saved to the running configuration.

You can change this default behavior by switching to the Transaction-Based Configuration mode. To switch to Transaction-Based Configuration mode, enter the `start transaction` command. When you switch to the Transaction-Based Configuration mode, you update the candidate configuration. Changes to the candidate configuration are not added to the running configuration until you commit them, which activates the configuration. The `start transaction` command applies only to the current session. Changing the configuration mode of the current session to the Transaction-Based Configuration mode does not affect the configuration mode of other CLI sessions.

- After you explicitly enter the `commit` command to save changes to the candidate configuration, the session switches back to the default behavior of automatically saving the configuration changes to the running configuration.
- When a session terminates while in the Transaction-Based Configuration mode, and you have not entered the `commit` command, the changes are maintained in the candidate configuration. You can start a new Transaction-Based Configuration mode session and continue with the remaining configuration changes.
- All sessions in Transaction-Based Configuration mode update the same candidate configuration. When you enter the `commit` command on any session in Transaction-Based Configuration mode or you make configuration changes on any session in Non-Transaction-Based mode, you also commit the changes made to the candidate configuration in all other sessions running in the transaction-based configuration mode. This implies that inconsistent configuration changes may be applied to the running configuration. Dell EMC recommends only making configuration changes on a single CLI session at a time.
- When you enter the `lock` command in a CLI session, configuration changes are disabled on all other sessions, whether they are in Transaction-Based Configuration mode or Non-Transaction-Based Configuration mode. For more information, see Candidate configuration.

CLI command hierarchy

CLI commands are organized in a hierarchy. Commands that perform a similar function are grouped together under the same level of hierarchy. For example, all commands that display information about the system and the system software are grouped under the `show system` command, and all commands that display information about the routing table are grouped under the `show ip route` command.

CLI command categories

There are several broad groups of CLI commands available:

copy	Copies files from one location on a device to another, from a device to a remote system, or from a remote system to a device.
configure	Enters CONFIGURATION mode to configure routing protocols, interfaces, network management, and user access.

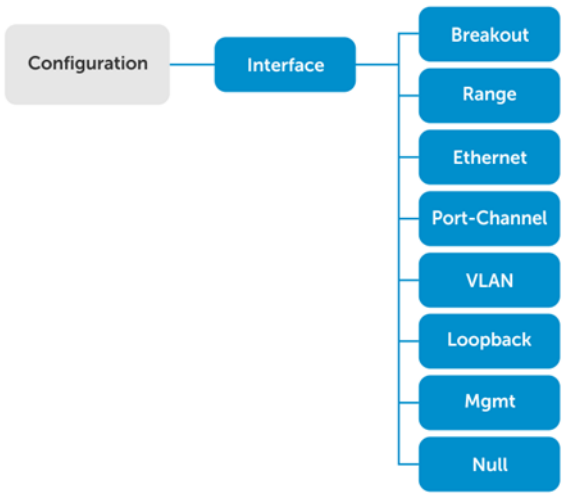
exit	Moves up one command mode. To go directly to EXEC mode, use the `end` command.

CONFIGURATION Mode

When you initially log in to OS10, you are placed in EXEC mode. To access CONFIGURATION mode, enter the `configure terminal` command. Use CONFIGURATION mode to manage interfaces, protocols, and features.

Interface mode is a sub-mode of CONFIGURATION mode. Interface mode is where you configure Layer 2 (L2) and Layer 3 (L3) protocols, and IPv4 and IPv6 services specific to an interface:

- Physical interfaces include the Management interface and Ethernet ports
- Logical interfaces include Loopback, port-channel, and virtual local area networks (VLANs)



From CONFIGURATION mode, you can also configure L2 and L3 protocols with a specific protocol-configuration mode, such as Spanning-Tree Protocol (STP) or Border Gateway Protocol (BGP).

Command help

To view a list of valid commands for any CLI mode, enter `?`.

1 Enter `?` to view the commands available in EXEC mode.

```
OS10# ?
  alarm                    Alarm commands
  alias                    Set alias for a command
  batch                    Batch Mode
  boot                     Tell the system where to access the software image at bootup
  clear                    Clear command
  clock                    Configure the system clock
  commit                   Commit candidate configuration
  configure                Enter configuration mode
  copy                     Perform a file copy operation
  crypto                   Cryptography commands
  debug                    Debug command
  delete                   Perform a file delete operation on local file system
  dir                      Show the list of files for the specified system folder
  discard                  Discard candidate configuration
  exit                     Exit from the CLI
  generate                 Command to generate executed functionality
  help                     Display available commands
  image                    Image commands
```

```
  kill-session              Kill a CLISH session
  license                   License and digital fulfillment commands
  location-led              Set location LED
  lock                      Lock candidate configuration
  move                      Perform a file move/rename operation on local filesystem
  no                        No commands under exec mode
  ping                      ping -h shows help
  ping6                     ping6 -h shows help
  reload                    Reboot Networking Operating System
  show                      Show running system information
  start                     Activate transaction based configuration
  support-assist-activity   Support Assist related activity
  system                    System command
  terminal                  Set terminal settings
  traceroute                traceroute --help shows help
  unlock                    Unlock candidate configuration
  validate                  Validate candidate configuration
  write                     Copy from current system configuration
  ztd                       Cancel the current ZTD process.
```

2 Enter CONFIGURATION mode.

```
OS10# configure terminal
OS10(config)#
```

3 Enter ? to show the commands available in CONFIGURATION mode.

```
OS10(config)# ?
  aaa                       Configure AAA
  alias                     Set alias for a command
  banner                    Configure banners
  bfd                       Enable bfd globally
  class-map                 Configure class map
  clock                     Configure clock parameters
  control-plane             Control-plane configuration
  crypto                    Crypto commands
  dcbx                      DCBX commands
  default                   Configure default attributes
  dot1x                     Configure dot1x global information
  end                       Exit to the exec Mode
  errdisable                Configure errdisable parameters
  eula-consent              eula-consent configuration
  evpn                      Ethernet Virtual Private Network (EVPN)
  exec-timeout              Set timeout (in seconds) for all CLI sessions
  exit                      Exit from current mode
  fcoe                      Configure FCoE
  feature                   Enable feature
  hardware                  Hardware specific configurations
  hash-algorithm            Hash algorithm configurations
  help                      Display available commands
  host-description          Set the system host description
  hostname                  Set the system hostname
  interface                 Select an interface
  ip                        Global IP configuration subcommands
  ipv6                      Global ipv6 configuration
  iscsi                     Enable iscsi globally
  lacp                      LACP commands
  line                      Configure a terminal line
  link-bundle-utilization   Configure link bundle utilization trigger threshold
  lldp                      Configure LLDP parameters
  load-balancing            Load balancing configurations
  logging                   Logging commands
  login                     Configure login parameters
  mac                       MAC config commands
  management                management interface commands
  monitor                   Create a session for monitoring traffic
  no                        To delete / disable commands in config mode
  ntp                       Configure NTP
  nve                       Create a Network Virtualization Edge (NVE) instance
  openflow                  Configure OpenFlow
```

```
  password-attributes      Configure the password attributes
  policy-map               Configure policy map
  qos-map                  Configure QoS map
  radius-server            Specify radius server host and configure its communication
parameters
  rest                     Configure rest interface
  route-map                Creates route-map
  router                   Enable a routing process
  scale-profile            Configure scale profile
  sflow                    Configure sflow parameters
  snmp-server              Configure SNMP server
  spanning-tree            Spanning Tree Subsystem
  support-assist           Support Assist feature configuration
  system                   System configuration
  tacacs-server            Specify TACACS+ server host and configure its communication
parameters
  track                    Configure object tracking
  trust                    Configure trust
  uplink-state-group       Create uplink state group
  username                 Create or modify users
  userrole                 Create custom user role
  virtual-network          Create a Virtual Network
  vlt-domain               VLT domain configurations
  vrrp                     Configure VRRP global attributes
  wred                     Configure WRED profile
```

Check device status

Use `show` commands to check the status of a device and monitor activities.

- Enter `show ?` from EXEC mode to view a list of commands to monitor a device.

```
OS10# show ?
  acl-table-usage          Show ACL table utilization
  alarms                   Display all current alarm situation in the system
  alias                    Show list of aliases
  bfd                      Show bfd session commands
  boot                     Show boot information
  candidate-configuration  Current candidate configuration
  class-map                Show QoS class-map configuration
  clock                    Show the system date and time
  command-history          Show command history of the current user
  control-plane            Display control-plane related informations
  copy-file                Show file copy operation information
  crypto                   Display cryptographic information
  diag                     Show diagnostic information for port adapters/modules
  diff                     Display differences between two configuration set
  discovered-expanders     discovered expanders info
  dot1x                    Show dot1x information
  environment              Show the environmental information of the system
  errdisable               Show errdisable information
  eula-consent             Shows eula-consent for various modules
  evpn                     Show Ethernet Virtual Private Network
  exec-timeout             Show the timeout value of CLI session (in seconds)
  fcoe                     show fcoe
  file                     Display file content in specified location
  fips                     Show fips mode status
  hardware                 Show hardware information
  hash-algorithm           Show hash algorithm information
  hosts                    show information about DNS
  image                    Show image information
  interface                Interface status and configuration
  inventory                Show the system inventory information
  ip                       show IP commands
  ipv6                     Display IPv6 neighbor information
  iscsi                    Show iscsi
```

```
  lacp                      Show LACP information
  license                   Show license and digital fulfillment related information
  link-bundle-utilization   Display the link-bundle utilization for the interfaces in the
bundle
  lldp                      Show lldp
  load-balance              Show global traffic load-balance configuration
  logging                   Show logging messages
  login                     Show login parameters
  mac                       MAC forwarding table
  monitor                   Show port monitoring sessions
  network-policy            Show network policy
  ntp                       NTP associations
  nve                       Display NVE related info
  parser-tree               Show parser tree
  policy-map                Show policy-map information
  port-channel              LAG status and configuration
  processes                 Show processes statistics
  qos                       Show ingress or egress QoS configuration
  queuing                   Show egress QoS counters
  route-map                 Show route map information
  running-configuration     Current operating configuration
  sessions                  Show active management sessions
  sflow                     Show sflow
  snmp                      Display all current snmp configuration
  spanning-tree             Show spanning tree information
  startup-configuration     Contents of startup configuration
  storm-control             Show storm control configuration
  support-assist            Show information about the support assist module
  switch-operating-mode     Switch operating mode
  system                    Show system status information
  tech-support              Collection of show commands
  terminal                  Show terminal configurations for this session
  trace                     Show trace messages
  track                     Show object tracking information
  uplink-state-group        Display the uplink state group configurations
  uptime                    Show the system uptime
  users                     Show the current list of users logged into the system , and show
the session id
  version                   Show the software version on the system
  virtual-network           Virtual-network info
  vlan                      Vlan status and configuration
  vlt                       Show VLT domain info
  vrrp                      VRRP group status
  ztd-status                Show ztd status
```

- Enter `show command-history` from EXEC mode to view trace messages for each executed command.

```
OS10# show command-history
    1    Thu Apr  20 19:44:38 UTC 2017  show vlan
    2    Thu Apr  20 19:47:01 UTC 2017  admin
    3    Thu Apr  20 19:47:01 UTC 2017  monitor hardware-components controllers view 0
    4    Thu Apr  20 19:47:03 UTC 2017  system general info system-version view
    5    Thu Apr  20 19:47:16 UTC 2017  admin
    6    Thu Apr  20 19:47:16 UTC 2017  terminal length 0
    7    Thu Apr  20 19:47:18 UTC 2017  terminal datadump
    8    Thu Apr  20 19:47:20 UTC 2017  %abc
    9    Thu Apr  20 19:47:22 UTC 2017  switchshow
   10    Thu Apr  20 19:47:24 UTC 2017  cmsh
   11    Thu Apr  20 19:47:26 UTC 2017  show version
   12    Thu Apr  20 19:47:28 UTC 2017  cmsh
   13    Thu Apr  20 19:47:30 UTC 2017  show version
   14    Thu Apr  20 19:47:32 UTC 2017  show system
   15    Fri Apr  21 12:35:31 UTC 2017  BIOS 3.20.0.3
```

- Enter `clear command-history` to clear the trace messages in `show command-history`.

```
OS10# clear command-history
```

- Check the `show command-history` to verify that the trace messages are cleared.
- Enter `show system` from EXEC mode to view the system status information.

```
OS10# show system
```

```
Node Id              : 1
MAC                  : 34:17:eb:f2:9a:c4
Number of MACs       : 256
Up Time              : 2 days 05:57:17

-- Unit 1 --
Status                       : up
System Identifier            : 1
Down Reason                  : unknown
Digital Optical Monitoring   : disable
System Location LED          : off
Required Type                : S4048
Current Type                 : S4048
Hardware Revision            : X01
Software Version             : 10.4.3.0
Physical Ports               : 48x10GbE, 6x40GbE
BIOS                            : 3.21.0.4
System CPLD                     : 15
Master CPLD                     : 12
Slave CPLD                      : 5

-- Power Supplies --
PSU-ID  Status      Type    AirFlow   Fan  Speed(rpm)  Status
-------------------------------------------------------------------
1       not-present
2       up          UNKNOWN NORMAL    1    10704       up


-- Fan Status --
FanTray  Status      AirFlow   Fan  Speed(rpm)  Status
-------------------------------------------------------------------
1        up          NORMAL    1    10108       up
                               2    10069       up

2        up          NORMAL    1    9954        up
                               2    10108       up

3        up          NORMAL    1    9867        up
                               2    9867        up
```

Candidate configuration

When you enter OS10 configuration commands in Transaction-Based Configuration mode, changes do not take effect immediately and are stored in the candidate configuration. The configuration changes become active only after you commit the changes using the `commit` command. Changes in the candidate configuration are validated and applied to the running configuration.

The candidate configuration allows you to avoid introducing errors during an OS10 configuration session. You can make changes and then check them before committing them to the active, running configuration on the network device.

To check differences between the running configuration and the candidate configuration, use the `show diff` command. After comparing the two, decide if you will commit the changes to the running configuration. To delete uncommitted changes, use the `discard` command.

- Enter `show ?` from EXEC mode to view a list of commands to monitor a device.

```
OS10# show ?
  acl-table-usage          Show ACL table utilization
  alarms                   Display all current alarm situation in the system
  alias                    Show list of aliases
  bfd                      Show bfd session commands
  boot                     Show boot information
  candidate-configuration  Current candidate configuration
  class-map                Show QoS class-map configuration
  clock                    Show the system date and time
  command-history          Show command history of the current user
  control-plane            Display control-plane related informations
```

```
  copy-file                 Show file copy operation information
  crypto                    Display cryptographic information
  diag                      Show diagnostic information for port adapters/modules
  diff                      Display differences between two configuration set
  discovered-expanders      discovered expanders info
  dot1x                     Show dot1x information
  environment               Show the environmental information of the system
  errdisable                Show errdisable information
  eula-consent              Shows eula-consent for various modules
  evpn                      Show Ethernet Virtual Private Network
  exec-timeout              Show the timeout value of CLI session (in seconds)
  fcoe                      show fcoe
  file                      Display file content in specified location
  fips                      Show fips mode status
  hardware                  Show hardware information
  hash-algorithm            Show hash algorithm information
  hosts                     show information about DNS
  image                     Show image information
  interface                 Interface status and configuration
  inventory                 Show the system inventory information
  ip                        show IP commands
  ipv6                      Display IPv6 neighbor information
  iscsi                     Show iscsi
  lacp                      Show LACP information
  license                   Show license and digital fulfillment related information
  link-bundle-utilization  Display the link-bundle utilization for the interfaces in the
bundle
  lldp                      Show lldp
  load-balance              Show global traffic load-balance configuration
  logging                   Show logging messages
  login                     Show login parameters
  mac                       MAC forwarding table
  monitor                   Show port monitoring sessions
  network-policy            Show network policy
  ntp                       NTP associations
  nve                       Display NVE related info
  parser-tree               Show parser tree
  policy-map                Show policy-map information
  port-channel              LAG status and configuration
  processes                 Show processes statistics
  qos                       Show ingress or egress QoS configuration
  queuing                   Show egress QoS counters
  route-map                 Show route map information
  running-configuration     Current operating configuration
  sessions                  Show active management sessions
  sflow                     Show sflow
  snmp                      Display all current snmp configuration
  spanning-tree             Show spanning tree information
  startup-configuration     Contents of startup configuration
  storm-control             Show storm control configuration
  support-assist            Show information about the support assist module
  switch-operating-mode     Switch operating mode
  system                    Show system status information
  tech-support              Collection of show commands
  terminal                  Show terminal configurations for this session
  trace                     Show trace messages
  track                     Show object tracking information
  uplink-state-group        Display the uplink state group configurations
  uptime                    Show the system uptime
  users                     Show the current list of users logged into the system , and show
the session id
  version                   Show the software version on the system
  virtual-network           Virtual-network info
  vlan                      Vlan status and configuration
  vlt                       Show VLT domain info
  vrrp                      VRRP group status
  ztd-status                Show ztd status
```

Compressed configuration

OS10 offers the `show candidate-configuration compressed` and `show running-configuration compressed` commands that display interface-related configuration in a compressed manner. These commands group similar looking configuration. The compression is done only for interface-related configuration (VLAN and physical interfaces).

View compressed candidate configuration

```
OS10# show candidate-configuration compressed
interface breakout 1/1/1 map 40g-1x
interface breakout 1/1/2 map 40g-1x
interface breakout 1/1/3 map 40g-1x
interface breakout 1/1/4 map 40g-1x
interface breakout 1/1/5 map 40g-1x
interface breakout 1/1/6 map 40g-1x
interface breakout 1/1/7 map 40g-1x
interface breakout 1/1/8 map 40g-1x
interface breakout 1/1/9 map 40g-1x
interface breakout 1/1/10 map 40g-1x
interface breakout 1/1/11 map 40g-1x
interface breakout 1/1/12 map 40g-1x
interface breakout 1/1/13 map 40g-1x
interface breakout 1/1/14 map 40g-1x
interface breakout 1/1/15 map 40g-1x
interface breakout 1/1/16 map 40g-1x
interface breakout 1/1/17 map 40g-1x
interface breakout 1/1/18 map 40g-1x
interface breakout 1/1/19 map 40g-1x
interface breakout 1/1/20 map 40g-1x
interface breakout 1/1/21 map 40g-1x
interface breakout 1/1/22 map 40g-1x
interface breakout 1/1/23 map 40g-1x
interface breakout 1/1/24 map 40g-1x
interface breakout 1/1/25 map 40g-1x
interface breakout 1/1/26 map 40g-1x
interface breakout 1/1/27 map 40g-1x
interface breakout 1/1/28 map 40g-1x
interface breakout 1/1/29 map 40g-1x
interface breakout 1/1/30 map 40g-1x
interface breakout 1/1/31 map 40g-1x
interface breakout 1/1/32 map 40g-1x
ipv6 forwarding enable
username admin password $6$q9QBeYjZ$jfxzVqGhkxX3smxJSH9DDz7/3OJc6m5wjF8nnLD7/VKx8SloIhp4NoGZs0I/
UNwh8WVuxwfd9q4pWIgNs5BKH. role sysadmin
aaa authentication local
snmp-server contact http://www.dell.com/support
!
interface range ethernet 1/1/1-1/1/32
 switchport access vlan 1
 no shutdown
!
interface vlan 1
 no shutdown
!
interface mgmt1/1/1
 ip address dhcp
 no shutdown
 ipv6 enable
 ipv6 address autoconfig
!
support-assist
!
policy-map type application policy-iscsi
!
class-map type application class-iscsi
```

View compressed running configuration

```
OS10# show running-configuration compressed
interface breakout 1/1/1 map 40g-1x
interface breakout 1/1/2 map 40g-1x
interface breakout 1/1/3 map 40g-1x
interface breakout 1/1/4 map 40g-1x
interface breakout 1/1/5 map 40g-1x
interface breakout 1/1/6 map 40g-1x
interface breakout 1/1/7 map 40g-1x
interface breakout 1/1/8 map 40g-1x
interface breakout 1/1/9 map 40g-1x
interface breakout 1/1/10 map 40g-1x
interface breakout 1/1/11 map 40g-1x
interface breakout 1/1/12 map 40g-1x
interface breakout 1/1/13 map 40g-1x
interface breakout 1/1/14 map 40g-1x
interface breakout 1/1/15 map 40g-1x
interface breakout 1/1/16 map 40g-1x
interface breakout 1/1/17 map 40g-1x
interface breakout 1/1/18 map 40g-1x
interface breakout 1/1/19 map 40g-1x
interface breakout 1/1/20 map 40g-1x
interface breakout 1/1/21 map 40g-1x
interface breakout 1/1/22 map 40g-1x
interface breakout 1/1/23 map 40g-1x
interface breakout 1/1/24 map 40g-1x
interface breakout 1/1/25 map 40g-1x
interface breakout 1/1/26 map 40g-1x
interface breakout 1/1/27 map 40g-1x
interface breakout 1/1/28 map 40g-1x
interface breakout 1/1/29 map 40g-1x
interface breakout 1/1/30 map 40g-1x
interface breakout 1/1/31 map 40g-1x
interface breakout 1/1/32 map 40g-1x
ipv6 forwarding enable
username admin password $6$q9QBeYjZ$jfxzVqGhkxX3smxJSH9DDz7/3OJc6m5wjF8nnLD7/VKx8SloIhp4NoGZs0I/
UNwh8WVuxwfd9q4pWIgNs5BKH. role sysadmin
aaa authentication local
snmp-server contact http://www.dell.com/support
!
interface range ethernet 1/1/1-1/1/32
 switchport access vlan 1
 no shutdown
!
interface vlan 1
 no shutdown
!
interface mgmt1/1/1
 ip address dhcp
 no shutdown
 ipv6 enable
 ipv6 address autoconfig
!
support-assist
!
policy-map type application policy-iscsi
!
class-map type application class-iscsi
```

Show difference between candidate and running configurations

```
OS10# show diff candidate-configuration running-configuration
OS10#
```

NOTE: If the `show` command does not return output, the `candidate-configuration` and `running-configuration` files match.

Prevent configuration changes

You can prevent configuration changes on sessions other than the current CLI session using the `lock` command. To respectively prevent and allow configuration changes on other sessions, use the `lock` and `unlock` commands in EXEC mode. When you enter the `lock` command on a CLI session, users cannot make configuration changes across any other active CLI sessions. When you close the CLI session where you entered the `lock` command, configuration changes are automatically allowed on all other sessions.

Lock configuration changes

```
OS10# lock
```

Unlock configuration changes

```
OS10# unlock
```

Change to transaction-based configuration mode

To change to Transaction-Based Configuration mode for a session, enter the `start transaction` command.

1 Change to Transaction-Based Configuration mode in EXEC mode.

```
start transaction
```

2 Enable, for example, an interface from INTERFACE mode.

```
interface ethernet 1/1/1
no shutdown
```

3 Save the configuration.

```
do commit
```

> **NOTE: After you enter the `do commit` command, the current session switches back to the default behavior of committing all configuration changes automatically.**

Save configuration changes manually

```
OS10# start transaction
OS10# configure terminal
OS10(config)#
OS10(config)# interface ethernet 1/1/1
OS10(config-if-eth1/1/1)# no shutdown
OS10(config-if-eth1/1/1)# do commit
```

Copy running configuration

The running configuration contains the current OS10 system configuration and consists of a series of OS10 commands. Copy the running configuration to a remote server or local directory as a backup or for viewing and editing. The running configuration is copied as a text file, which you can view and edit with a text editor. To copy the running configuration to the startup configuration file, enter the `copy running-configuration startup-configuration` command.

Copy running configuration to local directory or remote server

```
OS10# copy running-configuration {config://filepath | home://filepath |
ftp://userid:passwd@hostip/filepath | scp://userid:passwd@hostip/filepath |
sftp://userid:passwd@hostip/filepath | tftp://hostip/filepath}
```

```
OS10# copy running-configuration scp://root:calvin@10.11.63.120/tmp/qaz.txt
```

Copy file to running configuration

To apply a set of commands to the current running configuration and execute them immediately, copy a text file from a remote server or local directory. The copied commands do not replace the existing commands. If the copied command fails, the successful copied commands before the failure is maintained.

```
OS10# copy {config://filepath | home://filepath |
ftp://userid:passwd@hostip/filepath | scp://userid:passwd@hostip/filepath |
sftp://userid:passwd@hostip/filepath | tftp://hostip/filepath | http://userid@hostip/filepath}
running-configuration
```

```
OS10# copy scp://root:calvin@10.11.63.120/tmp/qaz.txt running-configuration
```

Restore startup configuration

The startup configuration file, `startup.xml`, is stored in the `config` system folder. To create a backup version, copy the startup configuration to a remote server or the local `config:` or `home:` directories.

To restore a backup configuration, copy a local or remote file to the startup configuration and reload the switch. After downloading a backup configuration, enter the `reload` command, otherwise the configuration does not take effect until you reboot.

Copy file to startup configuration

```
OS10# copy {config://filepath | home://filepath |
ftp://userid:passwd@hostip/filepath | scp://userid:passwd@hostip/filepath |
sftp://userid:passwd@hostip/filepath | tftp://hostip/filepath} config://startup.xml
```

Back up startup file

```
OS10# copy config://startup.xml config://backup-9-28.xml
```

Restore startup file from backup

```
OS10# copy config://backup-9-28.xml config://startup.xml
OS10# reload
```

Back up startup file to server

```
OS10# copy config://startup.xml scp://userid:password@hostip/backup-9-28.xml
```

Restore startup file from server

```
OS10# copy scp://admin:admin@hostip/backup-9-28.xml config://startup.xml
OS10# reload
```

Reload system image

Reboot the system manually using the `reload` command in EXEC mode. You are prompted to confirm the operation.

```
OS10# reload

System configuration has been modified. Save? [yes/no]:yes

Saving system configuration

Proceed to reboot the system? [confirm yes/no]:yes
```

To configure the OS10 image loaded at the next system boot, enter the `boot system` command in EXEC mode.

```
boot system {active | standby}
```

- Enter `active` to load the primary OS10 image stored in the A partition.
- Enter `standby` to load the secondary OS10 image stored in the B partition.

Set next boot image

```
OS10# boot system standby
OS10# show boot
Current system image information:
====================================
Type       Boot Type   Active          Standby         Next-Boot
------------------------------------------------------------------------
Node-id 1  Flash Boot  [A] 10.2.9999E  [B] 10.2.9999E  [B] standby
```

Filter show commands

You can filter `show` command output to view specific information, or start the command output at the first instance of a regular expression or phrase.

display-xml	Displays in XML format.
except	Shows only text that does not match a pattern
find	Searches for the first occurrence of a pattern and displays all the subsequent configurations
grep	Shows only text that matches a pattern
no-more	Does not paginate output
save	Saves the output to a file

Display all output

```
OS10# show running-configuration | no-more
```

Alias command

The `alias` command allows you to create shortcuts for commonly used or long commands. You can also execute long commands along with their parameters.

The `alias` supports the following modes:

- Persistent mode — The alias is persistent and is used in other sessions. The aliases created in Configuration mode are persistent.
- Non-persistent mode — The alias is used only within the current session. After you close the session, the alias is removed from the switch. The aliases created in Exec mode are non-persistent.

NOTE: You cannot use existing keywords, parameters, and short form of keywords as alias names, nor can you create a shortcut for the `alias` command. The alias name is case-sensitive and can have a maximum of 20 characters.

- Create an alias in EXEC or CONFIGURATION mode — EXEC mode for non-persistent and CONFIGURATION mode for persistent aliases. The alias value is the actual command where you use `$n` to enter the input parameters. You can substitute `$n` with either numbers ranging from 1 to 9 or with an asterisk (*) and enter the parameters while executing the commands using the alias. Use asterisk (*) to represent any number of parameters. The maximum number of input parameters is 9.

  ```
  alias alias-name alias-value
  ```
- Execute the commands using the alias in the respective modes.
- View the current aliases.

  ```
  show alias [brief | detail]
  ```
- Use the `no` form of the command to delete an alias.

  ```
  no alias alias-name
  ```

Create alias

```
OS10# alias showint "show interface $*"
OS10(config)# alias goint "interface ethernet $1"
```

View alias output for `showint`

```
OS10# showint status
--------------------------------------------------------------------------------
Port            Description     Status   Speed    Duplex   Mode Vlan Tagged-Vlans
--------------------------------------------------------------------------------
Eth 1/1/1                       up       40G              A    1    -
Eth 1/1/2                       up       40G              A    1    -
Eth 1/1/3                       up       40G              A    1    -
Eth 1/1/4                       up       40G              A    1    -
Eth 1/1/5                       up       40G              A    1    -
Eth 1/1/6                       up       40G              A    1    -
Eth 1/1/7                       up       40G              A    1    -
Eth 1/1/8                       up       40G              A    1    -
Eth 1/1/9                       up       40G              A    1    -
Eth 1/1/10                      up       40G              A    1    -
Eth 1/1/11                      up       40G              A    1    -
Eth 1/1/12                      up       40G              A    1    -
Eth 1/1/13                      up       40G              A    1    -
Eth 1/1/14                      up       40G              A    1    -
Eth 1/1/15                      up       40G              A    1    -
Eth 1/1/16                      up       40G              A    1    -
Eth 1/1/17                      up       40G              A    1    -
Eth 1/1/18                      up       40G              A    1    -
Eth 1/1/19                      up       40G              A    1    -
Eth 1/1/20                      up       40G              A    1    -
Eth 1/1/21                      up       40G              A    1    -
Eth 1/1/22                      up       40G              A    1    -
Eth 1/1/23                      up       40G              A    1    -
Eth 1/1/24                      up       40G              A    1    -
Eth 1/1/25                      up       40G              A    1    -
Eth 1/1/26                      up       40G              A    1    -
Eth 1/1/27                      up       40G              A    1    -
Eth 1/1/28                      up       40G              A    1    -
Eth 1/1/29                      up       40G              A    1    -
Eth 1/1/30                      up       40G              A    1    -
Eth 1/1/31                      up       40G              A    1    -
Eth 1/1/32                      up       40G              A    1    -
--------------------------------------------------------------------------------
```

View alias output for `goint`

```
OS10(config)# goint 1/1/1
OS10(conf-if-eth1/1/1)#
```

View alias information

```
OS10# show alias
Name                Type
----                ----
govlt               Config
goint               Config
shconfig            Local
showint             Local
shver               Local

Number of config aliases : 2
Number of local aliases : 3
```

View alias information brief. Displays the first 10 characters of the alias value.

```
OS10# show alias brief
Name                Type        Value
----                ----        -----
```

```
govlt                Config      "vlt-domain..."
goint                Config      "interface ..."
shconfig             Local       "show runni..."
showint              Local       "show inter..."
shver                Local       "show versi..."

Number of config aliases : 2
Number of local aliases : 3
```

View alias information in detail. Displays the entire alias value.

```
OS10# show alias detail
Name                 Type        Value
----                 ----        -----
govlt                Config      "vlt-domain $1"
goint                Config      "interface ethernet $1"
shconfig             Local       "show running-configuration"
showint              Local       "show interface $*"
shver                Local       "show version"

Number of config aliases : 2
Number of local aliases : 3
```

Delete alias

```
OS10#  no alias showint
OS10(config)# no alias goint
```

Multi-line alias

You can create a multi-line alias where you save a series of multiple commands in an alias. Multi-line alias is supported only in the Configuration mode.

You cannot use the exiting CLI keywords as alias names. The alias name is case-sensitive and can have a maximum of 20 characters.

- Create a multi-line alias in the CONFIGURATION mode. The switch enters the ALIAS mode.

  ```
  alias alias-name
  ```

- Enter the commands to be executed prefixed by the `line n` command in ALIAS mode. Enter the commands in double quotes and use `$n` to enter input parameters. You can substitute `$n` with either numbers ranging from 1 to 9 or with an asterisk (*) and enter the parameters while executing the commands using the alias. When you are using asterisk (*), you can use all the input parameters. The maximum number of input parameters is 9.

  ```
  line nn command
  ```

- (Optional) You can enter the default values to use for the parameters defined as `$n` in ALIAS mode.

  ```
  default n input-value
  ```

- (Optional) Enter a description for the multi-line alias in ALIAS mode.

  ```
  description string
  ```

- Use the `no` form of the command to delete an alias in CONFIGURATION mode.

  ```
  no alias alias-name
  ```

You can modify an existing multi-line alias by entering the corresponding ALIAS mode.

Create multi-line alias

```
OS10(config)# alias mTest
OS10(config-alias-mTest)# line 1 "interface $1 $2"
OS10(config-alias-mTest)# line 2 "no shutdown"
OS10(config-alias-mTest)# line 3 "show configuration"
OS10(config-alias-mTest)# default 1 "ethernet"
OS10(config-alias-mTest)# default 2 "1/1/1"
OS10(config-alias-mTest)# description InterfaceDetails
```

View alias output for `mTest` with default values

```
OS10(config)# mTest
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# show configuration
!
interface ethernet1/1/1
 no shutdown
 switchport access vlan 1
```

View alias output for `mTest` with different values

```
OS10(config)# mTest ethernet 1/1/10
OS10(config)# interface ethernet 1/1/10
OS10(conf-if-eth1/1/10)# no shutdown
OS10(conf-if-eth1/1/10)# show configuration
!
interface ethernet1/1/10
 no shutdown
 switchport access vlan 1
```

Modify existing multi-line alias

```
OS10(config)# alias mTest
OS10(config-alias-mTest)# line 4 "exit"
```

View the commands saved in the multi-line alias

```
OS10(config-alias-mTest)# show configuration
!
alias mTest
 description InterfaceDetails
 default 1 ethernet
 default 2 1/1/1
 line 1 "interface $1 $2"
 line 2 "no shutdown"
 line 3 "show configuration"
 line 4 exit
```

View alias information

```
OS10# show alias
Name                 Type
----                 ----
mTest                Config

Number of config aliases : 1
Number of local aliases : 0
```

View alias information brief. Displays the first 10 characters of each line of each alias.

```
OS10# show alias brief
Name                 Type        Value
----                 ----        -----
mTest                Config      line 1 "interface ..."
                                 line 2 "no shutdow..."
                                 line 3 "show confi..."
                                 default 1 "ethernet"
                                 default 2 "1/1/1"

Number of config aliases : 1
Number of local aliases : 0
```

View alias detail. Displays the entire alias value.

```
OS10# show alias detail
Name                 Type        Value
```

```
----                    ----          -----
mTest                   Config        line 1 "interface $1 $2"
                                      line 2 "no shutdown"
                                      line 3 "show configuration"
                                      default 1 "ethernet"
                                      default 2 "1/1/1"

Number of config aliases : 1
Number of local aliases : 0
```

Delete alias

```
OS10(config)# no alias mTest
```

Batch mode

Create and run a batch file to execute a sequence of multiple commands. A batch file is an unformatted text file that contains two or more commands. Store the batch file in the home directory.

Use vi or any other editor to create the batch file, then use the `batch` command to execute the file. To execute a series of commands in batch mode (non-interactive processing), use the `batch` command. OS10 automatically commits all commands in a batch file — you do not have to enter the `commit` command.

If a command in the batch file fails, batch operation stops at that command. The remaining commands are not executed.

- Create a batch file (for example, `b.cmd`) on a remote device by entering a series of commands.

  ```
  interface ethernet 1/1/1
  no shutdown
  no switchport
  ip address 172.17.4.1/24
  ```

- Copy the command file to the home directory on the switch.

  ```
  OS10# copy scp://os10user:os10passwd@10.11.222.1/home/os10/b.cmd home://b.cmd

  OS10# dir home

  Directory contents for folder: home
  Date (modified)         Size (bytes)   Name
  ---------------------   ------------   ------
  2017-02-15T19:25:35Z    77
  b.cmd
  ...
  ```

- Execute the batch file using the `batch /home/`*`username`*`/`*`filename`* command in EXEC mode.

  ```
  OS10# batch /home/admin/b.cmd
  Jun 26 18:29:12 OS10 dn_l3_core_services[723]: Node.1-Unit.1:PRI:notice [os10:trap],
  %Dell EMC (OS10) %log-notice:IP_ADDRESS_ADD: IP Address add is successful.
  IP 172.17.4.1/24 in VRF:default added successfully
  ```

- (Optional) Verify the new commands in the running configuration.

  ```
  OS10# show running-configuration interface ethernet 1/1/1
  !
  interface ethernet1/1/1
  no shutdown
  no switchport
  ip address 172.17.4.1/24
  ```

Linux shell commands

You can execute a single command, or a series of commands, using a batch file from the Linux shell.

- Use the `-c` option to run a single command.

```
admin@OS10:/opt/dell/os10/bin$ clish -c "show version"

New user admin logged in at session 10
Dell EMC Networking OS10 Enterprise
Copyright (c) 1999-2019 by Dell Inc. All Rights Reserved.
OS Version: 10.4.3.0
Build Version: 10.4.3.85
Build Time: 2019-02-18T17:06:10-0800
System Type: S4048-ON
Architecture: x86_64
Up Time: 2 days 05:58:01

User admin logged out at session 10
admin@OS10:/opt/dell/os10/bin$
```

- Use the `-B` option along with a batch file to execute a series of commands.

```
configure terminal
router bgp 100
neighbor 100.1.1.1
remote-as 104
no shutdown
```

Execute the batch file.

```
admin@OS10:/opt/dell/os10/bin$ clish -B ~/batch_cfg.txt

New user admin logged in at session 15
```

Verify the `BGP` configuration executed by the batch file.

```
admin@OS10:/opt/dell/os10/bin$ clish -c "show running-configuration bgp"

New user admin logged in at session 16
!
router bgp 100
 !
 neighbor 100.1.1.1
  remote-as 104
  no shutdown
admin@OS10:/opt/dell/os10/bin$

User admin logged out at session 16
```

SSH commands

You can execute commands remotely using a secure shell (SSH) session. This is supported only for `show` commands.

- Enter the `show` command along with `SSH`.

```
$ ssh admin@ip-address show-command
```

```
$ ssh admin@10.11.98.39 "show version"
admin@10.11.98.39's password:
Dell EMC Networking OS10 Enterprise
Copyright (c) 1999-2019 by Dell Inc. All Rights Reserved.
OS Version: 10.4.3.0
Build Version: 10.4.3.85
Build Time: 2019-02-18T17:06:10-0800
System Type: S4048-ON
Architecture: x86_64
Up Time: 2 days 05:58:01
```

OS9 environment commands

You can configure commands in an OS9 environment using the `feature config-os9-style` command.

- VLAN Interface mode
  - `tagged`
  - `no tagged`
  - `untagged`
  - `no untagged`
- Port-channel Interface mode:
  - `channel-member`
  - `no channel-member`
- Enable the feature to configure commands in an OS9 environment in CONFIGURATION mode.

```
OS10(config)# feature config-os9-style
OS10(config)# exit
OS10# show running-configuration compressed
interface breakout 1/1/28 map 10g-4x
feature config-os9-style
```

- After you enable this feature, you can use the OS9 format of commands only in the new session. This configuration does not take effect in the current session.

```
OS10(config)# interface vlan 11
OS10(conf-if-vl-11)# tagged ethernet 1/1/15

OS10(conf-if-vl-11)# show configuration
!
interface vlan11
 no shutdown
 tagged ethernet 1/1/15
```

Common commands

alias

Creates a command alias.

Syntax

`alias alias-name alias-value`

Parameters

- *alias-name* — Enter the name of the alias. A maximum of 20 characters.
- *alias-value* — Enter the command to be executed within double quotes. Enter the `$` followed by either numbers ranging from **1** to **9** or with an asterisk (*****) and enter the parameters while executing the commands using the alias. Use asterisk (*) to represent any number of parameters.

Default

Not configured

Command Mode

EXEC

CONFIGURATION

Usage Information Use this command to create a shortcut to long commands along with arguments. Use the numbers 1 to 9 along with the **$** to provide input parameters. The `no` version of this command deletes an alias.

Example In this example, when you enter `showint status`, note that the text on the CLI changes to `show interface status`. The alias changes to the actual command that you have specified in the alias definition.

```
OS10# alias showint "show interface $*"
OS10# showint status
--------------------------------------------------------------------------------
Port            Description     Status   Speed    Duplex   Mode Vlan Tagged-Vlans
--------------------------------------------------------------------------------
Eth 1/1/1                       up       40G              A    1    -
Eth 1/1/2                       up       40G              A    1    -
Eth 1/1/3                       up       40G              A    1    -
Eth 1/1/4                       up       40G              A    1    -
Eth 1/1/5                       up       40G              A    1    -
Eth 1/1/6                       up       40G              A    1    -
Eth 1/1/7                       up       40G              A    1    -
Eth 1/1/8                       up       40G              A    1    -
Eth 1/1/9                       up       40G              A    1    -
Eth 1/1/10                      up       40G              A    1    -
Eth 1/1/11                      up       40G              A    1    -
Eth 1/1/12                      up       40G              A    1    -
Eth 1/1/13                      up       40G              A    1    -
Eth 1/1/14                      up       40G              A    1    -
Eth 1/1/15                      up       40G              A    1    -
Eth 1/1/16                      up       40G              A    1    -
Eth 1/1/17                      up       40G              A    1    -
Eth 1/1/18                      up       40G              A    1    -
Eth 1/1/19                      up       40G              A    1    -
Eth 1/1/20                      up       40G              A    1    -
Eth 1/1/21                      up       40G              A    1    -
Eth 1/1/22                      up       40G              A    1    -
Eth 1/1/23                      up       40G              A    1    -
Eth 1/1/24                      up       40G              A    1    -
Eth 1/1/25                      up       40G              A    1    -
Eth 1/1/26                      up       40G              A    1    -
Eth 1/1/27                      up       40G              A    1    -
Eth 1/1/28                      up       40G              A    1    -
Eth 1/1/29                      up       40G              A    1    -
Eth 1/1/30                      up       40G              A    1    -
Eth 1/1/31                      up       40G              A    1    -
Eth 1/1/32                      up       40G              A    1    -
--------------------------------------------------------------------------------
```

In this example, when you enter `goint 1/1/1`, note that the text on the CLI changes to `interface ethernet 1/1/1`.

```
OS10(config)# alias goint "interface ethernet $1"
OS10(config)# goint 1/1/1
OS10(conf-if-eth1/1/1)#
```

Supported Releases 10.3.0E or later

alias (multi-line)

Creates a mulit-line command alias.

Syntax `alias` *alias-name*

Parameters *alias-name* — Enter the name of the multi-line alias. A maximum of up to 20 characters.

Default Not configured

Command Mode CONFIGURATION

Usage Information Use this command to save a series of multiple commands in an alias. The switch enters ALIAS mode when you create an alias. You can enter the series of commands to be executed using the `line` command. The `no` version of this command deletes an alias.

Example

```
OS10(config)# alias mTest
OS10(config-alias-mTest)# line 1 "interface $1 $2"
OS10(config-alias-mTest)# line 2 "no shutdown"
OS10(config-alias-mTest)# line 3 "show configuration"
```

Supported Releases 10.4.0E(R1) or later

batch

Executes a series of commands in a file in batch, non-interactive, processing.

Syntax `batch /home/username/filename`

Parameters

- *username* — Enter the user name that was used to copy the command file.
- *filename* — Enter the name of a batch command file.

Default Not configured

Command Mode EXEC

Usage Information Use this command to create a batch command file on a remote machine. Copy the command file to the home directory on your switch. This command executes commands in batch mode. OS10 automatically commits all commands in a batch file; you do not have to enter the `commit` command. To display the files stored in the home directory, enter `dir home`. Use the `dir home` command to view the files stored in the home directory.

Example

```
batch /home/admin/b.cmd
Jun 26 18:29:12 OS10 dn_l3_core_services[723]: Node.1-Unit.1:PRI:notice
[os10:trap],
%Dell EMC (OS10) %log-notice:IP_ADDRESS_ADD: IP Address add is successful.
IP 172.17.4.1/24 in VRF:default added successfully
```

Supported Releases 10.2.0E or later

boot

Configures which OS10 image to use the next time the system boots up.

Syntax `boot system [active | standby]`

Parameters

- `active` — Reset the running partition as the next boot partition.
- `standby` — Set the standby partition as the next boot partition.

Default Not configured

Command Mode	EXEC
Usage Information	Use this command to configure the location of the OS10 image used to reload the software at boot time. Use the `show boot` command to view the configured next boot image. This command is applied immediately.
Example	`OS10# boot system standby`
Supported Releases	10.2.0E or later

commit

Commits changes in the candidate configuration to the running configuration.

Syntax	`commit`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	Use this command to save changes to the running configuration. Use the `do commit` command to save changes in CONFIGURATION mode.
Example	`OS10# commit`
Example (configuration)	`OS10(config)# do commit`
Supported Releases	10.2.0E or later

configure

Enters CONFIGURATION mode from EXEC mode.

Syntax	`configure {terminal}`
Parameters	`terminal` — Enters CONFIGURATION mode from EXEC mode.
Default	Not configured
Command Mode	EXEC
Usage Information	Enter `conf t` for auto-completion.
Example	`OS10# configure terminal` `OS10(config)#`
Supported Releases	10.2.0E or later

copy

Copies the current running configuration to the startup configuration and transfers files between an OS10 switch and a remote device.

Syntax

```
copy [running-configuration startup-configuration | config://filepath |
coredump://filepath | ftp://filepath | home://filepath | scp://filepath |
sftp://filepath | supportbundle://filepath | tftp://filepath | usb://filepath]
```

Parameters

- `running-configuration startup-configuration` — (Optional) Copy the current running configuration file to the startup configuration file.
- `config://filepath` — (Optional) Copy from the configuration directory.
- `coredump://filepath` — (Optional) Copy from the coredump directory.
- `ftp://userid:passwd@hostip/filepath` — (Optional) Copy from a remote FTP server.
- `home://username/filepath` — (Optional) Copy from the home directory.
- `scp://userid:passwd@hostip/filepath` — (Optional) Copy from a remote SCP server.
- `sftp://userid:passwd@hostip/filepath` — (Optional) Copy from a remote SFTP server.
- `supportbundle://filepath` — (Optional) Copy from the support-bundle directory.
- `tftp://hostip/filepath` — (Optional) Copy from a remote TFTP server.
- `usb:filepath` — (Optional) Copy from a USB file system.

Default Not configured

Command Mode EXEC

Usage Information Use this command to save the running configuration to the startup configuration, transfer coredump files to a remote location, backup the startup configuration, retrieve a previously backed-up configuration, replace the startup configuration file, or transfer support bundles.

NOTE: Use the `image download` command to download the software image to the current active partition.

CAUTION: Please do not use copy commands to download the image to the switch, as it may install the image on the config partition resulting in loss of disk space for critical system applications and functions.

Example

```
OS10# dir coredump

Directory contents for folder: coredump
Date (modified)         Size (bytes)   Name
---------------------   ------------   ----------------------------------------
2017-02-15T19:05:41Z    12402278       core.netconfd-pro.
2017-02-15_19-05-09.gz

OS10# copy coredump://core.netconfd-pro.2017-02-15_19-05-09.gz scp://
os10user:os10passwd@10.11.222.1:/home/os10/core.netconfd-pro.2017-02
-15_19-05-09.gz
```

Example (copy startup configuration)

```
OS10# dir config

Directory contents for folder: config
Date (modified)         Size (bytes)   Name
---------------------   ------------   ------------
2017-02-15T20:38:12Z    54525
startup.xml
```

```
OS10# copy config://startup.xml scp://os10user:os10passwd@10.11.222.1:/home/
os10/backup.xml
```

Example (retrieve backed-up configuration)

```
OS10# copy scp://os10user:os10passwd@10.11.222.1:/home/os10/backup.xml home://
config.xml

OS10(conf-if-eth1/1/5)# dir home

Directory contents for folder: home
Date (modified)         Size (bytes)  Name
----------------------  ------------  -----------
…
2017-02-15T21:19:54Z    54525
config.xml
…
```

Example (replace startup configuration)

```
OS10# home://config.xml config://startup.xml
```

Supported Releases 10.2.0E or later

default (alias)

Configures default values for input parameters in a multi-line alias.

Syntax

```
default n value
```

Parameters

- *n* — Enter the number of the argument, from 1 to 9.
- *value* — Enter the value for the input parameter.

Default Not configured

Command Mode ALIAS

Usage Information To use special characters in the input parameter value, enclose the string in double quotes. The `no` version of this command removes the default value.

Example

```
OS10(config)# alias mTest
OS10(config-alias-mTest)# default 1 "ethernet 1/1/1"
```

Supported Releases 10.4.0E(R1) or later

delete

Removes or deletes the startup configuration file.

Syntax

```
delete [config://filepath | coredump://filepath | home://filepath | image://
filepath | startup-configuration | supportbundle://filepath | usb://filepath]
```

Parameters

- `config://`*filepath* — (Optional) Delete from the configuration directory.
- `coredump://`*filepath* — (Optional) Delete from the coredump directory.
- `home://`*filepath* — (Optional) Delete from the home directory.
- `image://`*filepath* — (Optional) Delete from the image directory.
- `startup-configuration` — (Optional) Delete the startup configuration.

- `supportbundle://filepath` — (Optional) Delete from the support-bundle directory.
- `usb://filepath` — (Optional) Delete from the USB file system.

Default Not configured

Command Mode EXEC

Usage Information Use this command to remove a regular file, software image, or startup configuration. Removing the startup configuration restores the system to the factory default. You need to reboot the switch using the `reload` command for the operation to take effect.

NOTE:

- Use caution when removing the startup configuration.
- When the config partition has low disk space, you will see a syslog message as below:

```
SYS_STAT_LOW_DISK_SPACE: Warning! Configuration directory has 0.0% free.
Please delete unnecessary files from home directory.
```

When you see such errors, please delete unwanted files from the home directory or you may encounter degraded system performance.

Example

```
OS10# delete startup-configuration
```

Supported Releases 10.2.0E or later

description (alias)

Configures a textual description for a multi-line alias.

Syntax `description string`

Parameters *string* — Enter a text string for a multi-line alias description.

Default Not configured

Command Mode ALIAS

Usage Information

- To use special characters as a part of the description string, enclose the string in double quotes.
- Spaces between characters are not preserved after entering this command unless you enclose the entire description in quotation marks, for example, "*text description*."
- Enter a text string after the `description` command to overwrite any previous text strings that you configured as the description.
- The `no` version of this command removes the description.

Example

```
OS10(config)# alias mTest
OS10(config-alias-mTest)# description "This alias configures interfaces"
```

Supported Releases 10.4.0E(R1) or later

dir

Displays files stored in available directories.

Syntax `dir {config | coredump | home | image | supportbundle | usb}`

Parameters

- `config` — (Optional) Folder containing configuration files.
- `coredump` — (Optional) Folder containing coredump files.
- `home` — (Optional) Folder containing files in user's home directory.
- `image` — (Optional) Folder containing image files.
- `supportbundle` — (Optional) Folder containing support bundle files.
- `usb` — (Optional) Folder containing files on a USB drive.

Default Not configured

Command Mode EXEC

Usage Information Use the `dir config` command to display configuration files. This command requires at least one parameter.

Example

```
OS10# dir
  config         Folder containing configuration files
  coredump       Folder containing coredump files
  home           Folder containing files in user's home directory
  image          Folder containing image files
  supportbundle  Folder containing support bundle files
```

Example (config)

```
OS10# dir config
Directory contents for folder: config
Date (modified)         Size (bytes)  Name
----------------------  ------------  -----------
2017-04-26T15:23:46Z    26704         startup.xml
```

Supported Releases 10.2.0E or later

discard

Discards changes made to the candidate configuration file.

Syntax `discard`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# discard
```

Supported Releases 10.2.0E or later

do

Executes most commands from all CONFIGURATION modes without returning to EXEC mode.

Syntax `do` *command*

Parameters *command* — Enter an EXEC-level command.

Default Not configured

Command Mode	INTERFACE
Usage Information	None

Example	`OS10(config)# interface ethernet 1/1/7`
	`OS10(conf-if-eth1/1/7)# no shutdown`
	`OS10(conf-if-eth1/1/7)# do show running-configuration`
	`...`
	`!`
	`interface ethernet1/1/7`
	`  no shutdown`
	`!`
	`...`

Supported Releases	10.2.0E or later

feature config-os9-style

Configure commands in an OS9 environment.

Syntax	`feature config-os9-style`
Parameters	None
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	After you enable the feature to configure commands in OS9 format, log out of the session. In the next session, you can configure the commands in OS9 format. This feature does not have an impact on the `show` commands. Use the `no` form of the command to disable the feature.

Example	`OS10(config)# feature config-os9-style`
	`OS10# show running-configuration compressed`
	`interface breakout 1/1/28 map 10g-4x`
	`feature config-os9-style`

Supported Releases	10.3.0E or later

exit

Returns to the next higher command mode.

Syntax	`exit`
Parameters	None
Default	Not configured
Command Mode	All
Usage Information	None
Example	`OS10(conf-if-eth1/1/1)# exit` `OS10(config)#`
Supported Releases	10.2.0E or later

license

Installs a license file from a local or remote location.

Syntax

```
license install [ftp: | http: | localfs: | scp: | sftp: | tftp: | usb:]
filepath
```

Parameters

- `ftp:` — (Optional) Install from the remote file system (`ftp://userid:passwd@hostip/filepath`).
- `http[s]:` — (Optional) Install from the remote file system (`http://hostip/filepath`).
- `http[s]:` — (Optional) Request from remote server (`http://hostip`).
- `localfs:` — (Optional) Install from the local file system (`localfs://filepath`).
- `scp:` — (Optional) Request from the remote file system (`scp://userid:passwd@hostip/filepath`).
- `sftp:` — (Optional) Request from the remote file system (`sftp://userid:passwd@hostip/filepath`).
- `tftp:` — (Optional) Request from the remote file system (`tftp://hostip/filepath`).
- `usb:` — (Optional) Request from the USB file system (`usb://filepath`).

Default Not configured

Command Mode EXEC

Usage Information Use this command to install the Enterprise Edition license file. For more information, see Download OS10 image and license. OS10 requires a perpetual license to run beyond the 120-day trial period. The license file is installed in the */mnt/license* directory.

Example

```
OS10# license install scp://user:userpwd/10.1.1.10/CFNNX42-NOSEnterprise-
License.lic
License installation success.
```

Supported Releases 10.3.0E or later

line (alias)

Configures the commands to be executed in a multi-line alias.

Syntax

```
line nn command
```

Parameters

- *nn* — Enter the line number, from 1 to 99. The commands are executed in the order of the line numbers.
- *command* — Enter the command to be executed enclosed in double quotes.

Default Not configured

Command Mode ALIAS

Usage Information The `no` version of this command removes the line number and the corresponding command from the multi-line alias.

Example

```
OS10(config)# alias mTest
OS10(config-alias-mTest)# line 1 "interface $1 $2"
OS10(config-alias-mTest)# line 2 "no shutdown"
OS10(config-alias-mTest)# line 3 "show configuration"
```

Supported Releases 10.4.0E(R1) or later

lock

Locks the candidate configuration and prevents any configuration changes on any other CLI sessions, either in Transaction or Non-Transaction-Based Configuration mode.

Syntax `lock`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information The `lock` command fails if there are uncommitted changes in the candidate configuration.

Example

```
OS10# lock
```

Supported Releases 10.2.0E or later

management route

Configures an IPv4/IPv6 static route the Management port uses. Repeat the command to configure multiple management routes.

Syntax `management route {`*`ipv4-address/mask`* `|` *`ipv6-address/prefix-length`*`} {`*`forwarding-router-address`* `| managementethernet}`

Parameters

- *`ipv4-address/mask`* — Enter an IPv4 network address in dotted-decimal format (A.B.C.D), then a subnet mask in prefix-length format (*/xx*).
- *`ipv6-address/prefix-length`* — Enter an IPv6 address in x:x:x:x::x format with the prefix length in /xxx format. The prefix range is /0 to /128.
- *`forwarding-router-address`* — Enter the next-hop IPv4/IPv6 address of a forwarding router (gateway) for network traffic from the Management port.
- `managementethernet` — Configure the Management port as the interface for the route; forces the route to be associated with the Management interface.

Default Not configured

Command Mode CONFIGURATION

Usage Information Management routes are separate from IP routes and are only used to manage the system through the Management port. To display the currently configured IPv4 and IPv6 management routes, enter the `show ip management-route` and `show ipv6 management-route` commands.

Example (IPv4)

```
OS10(config)# management route 10.10.20.0/24 10.1.1.1
OS10(config)# management route 172.16.0.0/16 managementethernet
```

Example (IPv6)

```
OS10(config)# management route 10::/64 10::1
```

Supported Releases 10.2.2E or later

move

Moves or renames a file on the configuration or home system directories.

Syntax `move [config: | home: | usb:]`

Parameters

- `config:` — Move from the configuration directory (`config://filepath`).
- `home:` — Move from the home directory (`home://filepath`).
- `usb:` — Move from the USB file system (`usb://filepath`).

Default Not configured

Command Mode EXEC

Usage Information Use the `dir config` command to view the directory contents.

Example

```
OS10# move config://startup.xml config://startup-backup.xml
```

Example (dir)

```
OS10# dir config

Directory contents for folder: config
Date (modified)         Size (bytes)  Name
----------------------  ------------  -----------
2017-04-26T15:23:46Z    26704         startup.xml
```

Supported Releases 10.2.0E or later

no

Disables or deletes commands in EXEC mode.

Syntax `no [alias | debug | support-assist-activity | terminal]`

Parameters

- `alias` — Remove an alias definition.
- `debug` — Disable debugging.
- `support-assist-activity` — SupportAssist-related activity.
- `terminal` — Reset terminal settings.

Default Not configured

Command Mode EXEC

Usage Information Use this command in EXEC mode to disable or remove a configuration. Use the `no ?` in CONFIGURATION mode to view available commands.

Example

```
OS10# no alias goint
```

Supported Releases 10.2.0E or later

reload

Reloads the software and reboots the ONIE-enabled device.

Syntax `reload`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information

NOTE: Use caution while using this command as it reloads the OS10 image and reboots the device.

Example

```
OS10# reload

Proceed to reboot the system? [confirm yes/no]:y
```

Supported Releases 10.2.0E or later

show alias

Displays configured alias commands available in both Persistent and Non-Persistent modes.

Syntax `show alias [brief | detail]`

Parameters

- `brief` — Displays brief information of the aliases.
- `detail` — Displays detailed information of the aliases.

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show alias
Name                 Type
----                 ----
govlt                Config
goint                Config
mTest                Config
shconfig             Local
showint              Local
shver                Local
Number of config aliases : 3
Number of local aliases : 3
```

Example (brief — displays the first 10 characters of the alias value))

```
OS10# show alias brief
Name                 Type        Value
----                 ----        -----
govlt                Config      "vlt-domain..."
goint                Config      "interface ..."
mTest                Config      line 1 "interface ..."
                                 line 2 "no shutdow..."
                                 line 3 "show confi..."
                                 default 1 "ethernet"
                                 default 2 "1/1/1"
shconfig             Local       "show runni..."
```

```
showint              Local       "show inter..."
shver                Local       "show versi..."

Number of config aliases : 3
Number of local aliases : 3
```

Example (detail — displays the entire alias value)

```
OS10# show alias detail
Name                 Type        Value
----                 ----        -----
govlt                Config      "vlt-domain $1"
goint                Config      "interface ethernet $1"
mTest                Config      line 1 "interface $1 $2"
                                 line 2 "no shutdown"
                                 line 3 "show configuration"
                                 default 1 "ethernet"
                                 default 2 "1/1/1"
shconfig             Local       "show running-configuration"
showint              Local       "show interface $*"
shver                Local       "show version"

Number of config aliases : 3
Number of local aliases : 3
```

Supported Releases 10.3.0E or later

show boot

Displays detailed information about the boot image.

Syntax `show boot [detail]`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information The `Next-Boot` field displays the partition that the next reload uses.

Example

```
OS10# show boot
Current system image information:
===================================
Type      Boot Type      Active                    Standby                  Next-Boot
-------------------------------------------------------------------------------------
Node-id 1 Flash Boot     [A] 10.4.3E             [B] 10.4.3E              [A] active
```

Example (detail)

```
OS10# show boot detail
Current system image information detail:
==========================================
Type:                      Node-id 1
Boot Type:                 Flash Boot
Active Partition:          A
Active SW Version:         10.4.3E
Active SW Build Version:   10.4.3E.85
Active Kernel Version:     Linux 4.9.110
Active Build Date/Time:    2019-02-18T09:06:10Z
Standby Partition:         B
Standby SW Version:        10.4.3E
Standby SW Build Version:  10.4.3E.80
Standby Build Date/Time:   2019-02-17T15:36:08Z
Next-Boot:                 active[A]
```

Supported Releases 10.2.0E or later

show candidate-configuration

Displays the current candidate configuration file.

Syntax

```
show candidate-configuration [aaa | access-list | as-path | bgp | class-map |
community-list | compressed | control-plane | dot1x | extcommunity-list |
interface | lacp | line | lldp | logging | monitor | ospf | ospfv3 | policy-map
| prefix-list | qos-map | radius-server | route-map | sflow | snmp | spanning-
tree | support-assist | system-qos | trust-map | users | vlt]
```

Parameters

- `aaa` — (Optional) Current candidate AAA configuration.
- `access-list` — (Optional) Current candidate access-list configuration.
- `as-path` — (Optional) Current candidate as-path configuration.
- `bgp` — (Optional) Current candidate BGP configuration.
- `class-map` — (Optional) Current candidate class-map configuration.
- `community-list` — (Optional) Current candidate community-list configuration.
- `compressed` — (Optional) Current candidate configuration in compressed format.
- `control-plane` — (Optional) Current candidate control-plane configuration.
- `dot1x` — (Optional) Current candidate dot1x configuration.
- `extcommunity-list` — (Optional) Current candidate extcommunity-list configuration.
- `interface` — (Optional) Current candidate interface configuration.
- `lacp` — (Optional) Current candidate LACP configuration.
- `lldp` — (Optional) Current candidate LLDP configuration.
- `logging` — (Optional) Current candidate logging configuration.
- `monitor` — (Optional) Current candidate monitor session configuration.
- `ospf` — (Optional) Current candidate OSPF configuration.
- `ospfv3` — (Optional) Current candidate OSPFv3 configuration.
- `policy-map` — (Optional) Current candidate policy-map configuration.
- `prefix-list` — (Optional) Current candidate prefix-list configuration.
- `qos-map` — (Optional) Current candidate qos-map configuration.
- `radius-server` — (Optional) Current candidate RADIUS server configuration.
- `route-map` — (Optional) Current candidate route-map configuration.
- `sflow` — (Optional) Current candidate sFlow configuration.
- `snmp` — (Optional) Current candidate SNMP configuration.
- `spanning-tree` — (Optional) Current candidate spanning-tree configuration.
- `support-assist` — (Optional) Current candidate support-assist configuration.
- `system-qos` — (Optional) Current candidate system-qos configuration.
- `trust-map` — (Optional) Current candidate trust-map configuration.
- `users` — (Optional) Current candidate users configuration.
- `vlt` — (Optional) Current candidate VLT domain configuration.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show candidate-configuration
! Version 10.2.9999E
! Last configuration change at Apr  11 10:36:43 2017
!
username admin password $6$q9QBeYjZ$jfxzVqGhkxX3smxJSH9DDz7/3OJc6m5wjF8nnLD7/
VKx8SloIhp4NoGZs0I/UNwh8WVuxwfd9q4pWIgNs5BKH.
aaa authentication local
snmp-server contact http://www.dell.com/support
snmp-server location "United States"
logging monitor disable
ip route 0.0.0.0/0 10.11.58.1
!
interface ethernet1/1/1
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/2
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/3
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/4
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/5
 switchport access vlan 1
 no shutdown
!
--more--
```

Example (compressed)

```
OS10# show candidate-configuration compressed
username admin password $6$q9QBeYjZ$jfxzVqGhkxX3smxJSH9DDz7/3OJc6m5wjF8nnLD7/
VKx8SloIhp4NoGZs0I/UNwh8WVuxwfd9q4pWIgNs5BKH.
aaa authentication local
snmp-server contact http://www.dell.com/support
snmp-server location "United States"
logging monitor disable
ip route 0.0.0.0/0 10.11.58.1
!
interface range ethernet 1/1/1-1/1/32
 switchport access vlan 1
 no shutdown
!
interface vlan 1
 no shutdown
!
interface mgmt1/1/1
 ip address 10.11.58.145/8
 no shutdown
 ipv6 enable
 ipv6 address autoconfig

!
support-assist
!
policy-map type application policy-iscsi
!
class-map type application class-iscsi
```

Supported Releases 10.2.0E or later

show environment

Displays information about environmental system components, such as temperature, fan, and voltage.

Syntax `show environment`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show environment

Unit    State              Temperature
-------------------------------------
1       up                 43

Thermal sensors
Unit   Sensor-Id        Sensor-name                                  Temperature
-------------------------------------------------------------------------------
1       1           CPU On-Board temp sensor                          32
1       2           Switch board  temp sensor                         28
1       3           System Inlet Ambient-1 temp sensor                27
1       4           System Inlet Ambient-2 temp sensor                25
1       5           System Inlet Ambient-3 temp sensor                26
1       6           Switch board 2 temp sensor                        31
1       7           Switch board 3 temp sensor                        41
1       8           NPU temp sensor                                   43
```

Supported Releases 10.2.0E or later

show inventory

Displays system inventory information.

Syntax `show inventory`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show inventory
Product                : S4048ON
Description            : S4048-ON 48x10GbE, 6x40GbE QSFP+ Interface Module
Software version       : 10.4.3.0
Product Base           :
Product Serial Number  :
Product Part Number    :

Unit Type                      Part Number  Rev  Piece Part ID             Svc Tag  Expr
-------------------------------------------------------------------------------------
* 1  S4048ON                   0J09D3       X01  TW-0J09D3-28298-49Q-0119  FFD7VS1  335
  1  S4048ON-PWR-2-UNKNOWN     0T9FNW       X01  TW-0T9FNW-28298-49Q-0041  AEIOU##  226
  1  S4048ON-FANTRAY-1         0MGDH8       X01  TW-0MGDH8-28298-49Q-0361  AEIOU##  226
```

```
  1   S4048ON-FANTRAY-2          0MGDH8        X01   TW-0MGDH8-28298-49Q-0360   AEIOU##   226
  1   S4048ON-FANTRAY-3          0MGDH8        X01   TW-0MGDH8-28298-49Q-0359   AEIOU##   226
```

Supported Releases 10.2.0E or later

show ip management-route

Displays the IPv4 routes used to access the Management port.

Syntax `show ip management-route [all | connected | summary]`

Parameters

- `all` — (Optional) Display the IPv4 routes that the Management port uses.
- `connected` — (Optional) Display only routes directly connected to the Management port.
- `summary` — (Optional) Display the number of active and non-active management routes and their remote destinations.
- `static` — (Optional) Display non-active management routes.

Default Not configured

Command Mode EXEC

Usage Information Use this command to view the IPv4 static and connected routes configured for the Management port. Use the `management route` command to configure an IPv4 or IPv6 management route.

Example

```
OS10# show ip management-route
 Destination        Gateway              State       Source
-------------------------------------------------------------------
192.168.10.0/24     managementethernet   Connected   Connected
```

Supported Releases 10.2.2E or later

show ipv6 management-route

Displays the IPv6 routes used to access the Management port.

Syntax `show ipv6 management-route [all | connected | summary]`

Parameters

- `all` — (Optional) Display the IPv6 routes that the Management port uses.
- `connected` — (Optional) Display only routes directly connected to the Management port.
- `summary` — (Optional) Display the number of active and non-active management routes and their remote destinations.
- `static` — (Optional) Display non-active Management routes.

Default Not configured

Command Mode EXEC

Usage Information Use this command to view the IPv6 static and connected routes configured for the Management port. Use the `management route` command to configure an IPv4 or IPv6 management route.

Example

```
OS10# show ipv6 management-route

Destination    Gateway                State
-----------    -------                -----
```

```
2001:34::0/64  ManagementEthernet 1/1  Connected
2001:68::0/64  2001:34::16             Active
```

Supported Releases 10.2.2E or later

show license status

Displays license status information.

Syntax `show license status`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information Use this command to view the `show license status` command to verify the current license for running OS10, its duration, and the service tag of the switch to which it is assigned.

Example

```
OS10# show license status

System Information
------------------------------------------------------------
Vendor Name          :   DELL
Product Name         :   S4048ON
Hardware Version     :   X01
Platform Name        :
PPID                 :   TW0J09D32829849Q0119
Service Tag          :   FFD7VS1
Product Base         :
Product Serial Number:
Product Part Number  :
License Details
----------------
Software        :        OS10-Enterprise
Version         :        10.4.3.0
License Type    :        PERPETUAL
License Duration:        Unlimited
License Status  :        Active
License location:        /mnt/license/FFD7VS1.lic
------------------------------------------------------------
```

Supported Releases 10.3.0E or later

show running-configuration

Displays the configuration currently running on the device.

Syntax `show running-configuration [aaa | access-list | as-path | bgp | class-map | community-list | compressed | control-plane | dot1x | extcommunity-list | interface [virtual-network vn-id] | lacp | line | lldp | logging | monitor | ospf | ospfv3 | policy-map | prefix-list | qos-map | radius-server | route-map | sflow | snmp | spanning-tree | support-assist | system-qos | trust-map | users | vlt]`

Parameters

- `aaa` — (Optional) Current operating AAA configuration.
- `access-list` — (Optional) Current operating access-list configuration.

- `as-path` — (Optional) Current operating as-path configuration.
- `bgp` — (Optional) Current operating BGP configuration.
- `class-map` — (Optional) Current operating class-map configuration.
- `community-list` — (Optional) Current operating community-list configuration.
- `compressed` — (Optional) Current operating configuration in compressed format.
- `control-plane` — (Optional) Current operating control-plane configuration.
- `dot1x` — (Optional) Current operating dot1x configuration.
- `extcommunity-list` — (Optional) Current operating extcommunity-list configuration.
- `interface` — (Optional) Current operating interface configuration.
  - `virtual-network` *vn-id* — (Optional) Current virtual network configuration.
- `lacp` — (Optional) Current operating LACP configuration.
- `lldp` — (Optional) Current operating LLDP configuration.
- `logging` — (Optional) Current operating logging configuration.
- `monitor` — (Optional) Current operating monitor session configuration.
- `ospf` — (Optional) Current operating OSPF configuration.
- `ospfv3` — (Optional) Current operating OSPFv3 configuration.
- `policy-map` — (Optional) Current operating policy-map configuration.
- `prefix-list` — (Optional) Current operating prefix-list configuration.
- `qos-map` — (Optional) Current operating qos-map configuration.
- `radius-server` — (Optional) Current operating radius-server configuration.
- `route-map` — (Optional) Current operating route-map configuration.
- `sflow` — (Optional) Current operating sFlow configuration.
- `snmp` — (Optional) Current operating SNMP configuration.
- `spanning-tree` — (Optional) Current operating spanning-tree configuration.
- `support-assist` — (Optional) Current operating support-assist configuration.
- `system-qos` — (Optional) Current operating system-qos configuration.
- `trust-map` — (Optional) Current operating trust-map configuration.
- `users` — (Optional) Current operating users configuration.
- `vlt` — (Optional) Current operating VLT domain configuration.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show running-configuration
! Version 10.2.9999E
! Last configuration change at Apr  11 01:25:02 2017
!
username admin password $6$q9QBeYjZ$jfxzVqGhkxX3smxJSH9DDz7/3OJc6m5wjF8nnLD7/
VKx8SloIhp4NoGZs0I/UNwh8WVuxwfd9q4pWIgNs5BKH.
aaa authentication local
snmp-server contact http://www.dell.com/support
snmp-server location "United States"
logging monitor disable
ip route 0.0.0.0/0 10.11.58.1
!
interface ethernet1/1/1
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/2
 switchport access vlan 1
```

```
 no shutdown
!
interface ethernet1/1/3
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/4
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/5
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/6
 switchport access vlan 1
 no shutdown
--more--
```

Example (compressed)

```
OS10# show running-configuration compressed
username admin password $6$q9QBeYjZ$jfxzVqGhkxX3smxJSH9DDz7/3OJc6m5wjF8nnLD7/
VKx8SloIhp4NoGZs0I/UNwh8WVuxwfd9q4pWIgNs5BKH.
aaa authentication local
snmp-server contact http://www.dell.com/support
snmp-server location "United States"
logging monitor disable
ip route 0.0.0.0/0 10.11.58.1
!
interface range ethernet 1/1/1-1/1/32
 switchport access vlan 1
 no shutdown
!
interface vlan 1
 no shutdown
!
interface mgmt1/1/1
 ip address 10.11.58.145/8
 no shutdown
 ipv6 enable
 ipv6 address autoconfig

!
support-assist
!
policy-map type application policy-iscsi
!
class-map type application class-iscsi
```

Supported Releases 10.2.0E or later

show startup-configuration

Displays the contents of the startup configuration file.

Syntax `show startup-configuration [compressed]`

Parameters `compressed` — (Optional) View a compressed version of the startup configuration file.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show startup-configuration
username admin password $6$q9QBeYjZ$jfxzVqGhkxX3smxJSH9DDz7/3OJc6m5wjF8nnLD7/
VKx8SloIhp4NoGZs0I/UNwh8WVuxwfd9q4pWIgNs5BKH.
aaa authentication local
snmp-server contact http://www.dell.com/support
snmp-server location "United States"
ip route 0.0.0.0/0 10.11.58.1
!
interface ethernet1/1/1
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/2
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/3
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/4
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/5
 switchport access vlan 1
 no shutdown
!
--more--
```

Example (compressed)

```
OS10# show startup-configuration compressed
username admin password $6$q9QBeYjZ$jfxzVqGhkxX3smxJSH9DDz7/3OJc6m5wjF8nnLD7/
VKx8SloIhp4NoGZs0I/UNwh8WVuxwfd9q4pWIgNs5BKH.
aaa authentication local
snmp-server contact http://www.dell.com/support
snmp-server location "United States"
ip route 0.0.0.0/0 10.11.58.1
!
interface range ethernet 1/1/1-1/1/32
 switchport access vlan 1
 no shutdown
!
interface vlan 1
 no shutdown
!
interface mgmt1/1/1
 ip address 10.11.58.145/8
 no shutdown
 ipv6 enable
 ipv6 address autoconfig

!
support-assist
!
policy-map type application policy-iscsi
!
class-map type application class-iscsi
```

Supported Releases 10.2.0E or later

show system

Displays system information.

Syntax `show system [brief | node-id]`

Parameters

- `brief` — View an abbreviated list of the system information.
- `node-id` — View the node ID number.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show system

Node Id              : 1
MAC                  : 34:17:eb:f2:9a:c4
Number of MACs       : 256
Up Time              : 2 days 05:57:17

-- Unit 1 --
Status                     : up
System Identifier          : 1
Down Reason                : unknown
Digital Optical Monitoring : disable
System Location LED        : off
Required Type              : S4048
Current Type               : S4048
Hardware Revision          : X01
Software Version           : 10.4.3.0
Physical Ports             : 48x10GbE, 6x40GbE
BIOS                          : 3.21.0.4
System CPLD                   : 15
Master CPLD                   : 12
Slave CPLD                    : 5

-- Power Supplies --
PSU-ID  Status      Type    AirFlow   Fan  Speed(rpm)  Status
-----------------------------------------------------------------
1       not-present
2       up          UNKNOWN NORMAL    1    10704       up


-- Fan Status --
FanTray  Status      AirFlow   Fan  Speed(rpm)  Status
-----------------------------------------------------------------
1        up          NORMAL    1    10108       up
                               2    10069       up

2        up          NORMAL    1    9954        up
                               2    10108       up

3        up          NORMAL    1    9867        up
                               2    9867        up
```

Example (node-id)

```
OS10# show system node-id 1 fanout-configured

Interface      Breakout capable     Breakout state
-----------------------------------------------------
Eth 1/1/1         No                BREAKOUT_1x1
Eth 1/1/2         No                BREAKOUT_1x1
```

```
Eth 1/1/3         No                BREAKOUT_1x1
Eth 1/1/4         No                BREAKOUT_1x1
Eth 1/1/5         No                BREAKOUT_1x1
Eth 1/1/6         No                BREAKOUT_1x1
Eth 1/1/7         No                BREAKOUT_1x1
Eth 1/1/8         No                BREAKOUT_1x1
Eth 1/1/9         No                BREAKOUT_1x1
Eth 1/1/10        No                BREAKOUT_1x1
Eth 1/1/11        No                BREAKOUT_1x1
Eth 1/1/12        No                BREAKOUT_1x1
Eth 1/1/13        No                BREAKOUT_1x1
Eth 1/1/14        No                BREAKOUT_1x1
Eth 1/1/15        No                BREAKOUT_1x1
Eth 1/1/16        No                BREAKOUT_1x1
Eth 1/1/17        No                BREAKOUT_1x1
Eth 1/1/18        No                BREAKOUT_1x1
Eth 1/1/19        No                BREAKOUT_1x1
Eth 1/1/20        No                BREAKOUT_1x1
Eth 1/1/21        No                BREAKOUT_1x1
Eth 1/1/22        No                BREAKOUT_1x1
Eth 1/1/23        No                BREAKOUT_1x1
Eth 1/1/24        No                BREAKOUT_1x1
Eth 1/1/25        No                BREAKOUT_1x1
Eth 1/1/26        No                BREAKOUT_1x1
Eth 1/1/27        No                BREAKOUT_1x1
Eth 1/1/28        No                BREAKOUT_1x1
Eth 1/1/29        No                BREAKOUT_1x1
Eth 1/1/30        No                BREAKOUT_1x1
Eth 1/1/31        No                BREAKOUT_1x1
Eth 1/1/32        No                BREAKOUT_1x1
Eth 1/1/33        No                BREAKOUT_1x1
Eth 1/1/34        No                BREAKOUT_1x1
Eth 1/1/35        No                BREAKOUT_1x1
Eth 1/1/36        No                BREAKOUT_1x1
Eth 1/1/37        No                BREAKOUT_1x1
Eth 1/1/38        No                BREAKOUT_1x1
Eth 1/1/39        No                BREAKOUT_1x1
Eth 1/1/40        No                BREAKOUT_1x1
Eth 1/1/41        No                BREAKOUT_1x1
Eth 1/1/42        No                BREAKOUT_1x1
Eth 1/1/43        No                BREAKOUT_1x1
Eth 1/1/44        No                BREAKOUT_1x1
Eth 1/1/45        No                BREAKOUT_1x1
Eth 1/1/46        No                BREAKOUT_1x1
Eth 1/1/47        No                BREAKOUT_1x1
Eth 1/1/48        No                BREAKOUT_1x1
Eth 1/1/49        Yes               BREAKOUT_1x1
Eth 1/1/50        Yes               BREAKOUT_1x1
Eth 1/1/51        Yes               BREAKOUT_1x1
Eth 1/1/52        Yes               BREAKOUT_1x1
Eth 1/1/53        Yes               BREAKOUT_1x1
Eth 1/1/54        Yes               BREAKOUT_1x1
```

Example (brief)

```
OS10# show system brief

Node Id           : 1
MAC               : 34:17:eb:f2:9a:c4

-- Unit --
Unit  Status      ReqType     CurType     Version
-----------------------------------------------------------------
1     up          S4048       S4048       10.4.3E

-- Power Supplies --
PSU-ID  Status      Type    AirFlow   Fan  Speed(rpm)  Status
-----------------------------------------------------------------
1       not-present
2       up          UNKNOWN NORMAL    1    10704       up
```

```
-- Fan Status --
FanTray  Status      AirFlow   Fan  Speed(rpm)  Status
--------------------------------------------------------------
1        up          NORMAL    1    9929        up
                               2    9980        up

2        up          NORMAL    1    10095       up
                               2    10082       up

3        up          NORMAL    1    9867        up
                               2    10173       up
```

Supported Releases 10.2.0E or later

show version

Displays software version information.

Syntax `show version`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
TR2# show version
Dell EMC Networking OS10 Enterprise
Copyright (c) 1999-2019 by Dell Inc. All Rights Reserved.
OS Version: 10.4.3.0
Build Version: 10.4.3.85
Build Time: 2019-02-18T17:06:10-0800
System Type: S4048-ON
Architecture: x86_64
Up Time: 2 days 05:58:01
```

Supported Releases 10.2.0E or later

start

Activates Transaction-Based Configuration mode for the active session.

Syntax `start` *`transaction`*

Parameters `transaction` - Enables the transaction-based configuration.

Default Not configured

Command Mode EXEC

Usage Information Use this command to save changes to the candidate configuration before applying configuration changes to the running configuration.

NOTE: Before you start a transaction, you must lock the session using the `lock` command in EXEC mode. Otherwise, the configuration changes from other sessions get committed.

Example	`OS10# start transaction`
Supported Releases	10.3.1E or later

system

Executes a Linux command from within OS10.

Syntax	`system` *command*
Parameters	*command* — Enter the Linux command to execute.
Default	Not configured
Command Mode	EXEC
Usage Information	None

Example	`OS10# system bash`
	`admin@OS10:~$ pwd`
	`/config/home/admin`
	`admin@OS10:~$ exit`
	`OS10#`

Supported Releases	10.2.0E or later

system-cli enable

Disables `system` command

Syntax	`system-cli enable`
Parameters	None
Default	None
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command enables OS10 `system` command.
Example	`OS10# configure terminal` `OS10(config)# system-cli enable`
Supported Releases	10.4.3.0 or later

system identifier

Sets a non-default unit ID in a non-stacking configuration.

Syntax	`system identifier` *system-identifier-ID*
Parameters	*system-identifier-ID* — Enter the system identifier ID, from 1 to 9.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The system ID displays in the stack LED on the switch front panel.

Example

```
OS10(config)# system identifier 1
```

Supported Releases 10.3.0E or later

terminal

Sets the number of lines to display on the terminal and enables logging.

Syntax

```
terminal {length lines | monitor}
```

Parameters

- `length lines` — Enter the number of lines to display on the terminal, from 0 to 512, default 24.
- `monitor` — Enables logging on the terminal.

Default 24 terminal lines

Command Mode EXEC

Usage Information Enter zero (`0`) for the terminal to display without pausing.

Example

```
OS10# terminal monitor
```

Supported Releases 10.2.0E or later

traceroute

Displays the routes that packets take to travel to an IP address.

Syntax

```
traceroute [vrf {management | vrf-name}] host [-46dFITnreAUDV] [-f first_ttl]
[-g gate,...] [-i device] [-m max_ttl] [-N squeries] [-p port] [-t tos] [-l
flow_label] [-w waittime] [-q nqueries] [-s src_addr] [-z sendwait] [--
fwmark=num] host [packetlen]
```

Parameters

- `vrf management`— (Optional) Traces the route to an IP address in the management VRF instance.
- `vrf vrf-name` — (Optional) Traces the route to an IP address in the specified VRF instance.
- `host` — Enter the host to trace packets from.
- `-i interface` — (Optional) Enter the IP address of the interface through which traceroute sends packets. By default, the interface is selected according to the routing table.
- `-m max_ttl` — (Optional) Enter the maximum number of hops, the maximum time-to-live value, that traceroute probes. The default is 30.
- `-p port` — (Optional) Enter a destination port:
  - For UDP tracing, enter the destination port base that traceroute uses. The destination port number is incremented by each probe.
  - For Internet Control Message Protocol (ICMP) tracing, enter the initial ICMP sequence value, incremented by each probe.
  - For TCP tracing, enter the constant destination port to connect.
  - `-P protocol` — (Optional) Use a raw packet of the specified protocol for traceroute. The default protocol is 253 (RFC 3692).
  - `-s source_address` — (Optional) Enter an alternative source address of one of the interfaces. By default, the address of the outgoing interface is used.
  - `-q nqueries` — (Optional) Enter the number of probe packets per hop. The default is 3.

- `-N squeries` — (Optional) Enter the number of probe packets sent out simultaneously to accelerate traceroute. The default is 16.
- `-t tos` — (Optional) For IPv4, enter the type of service (ToS) and precedence values to use. `16` sets a low delay; `8` sets a high throughput.
- `-UL` — (Optional) Use UDPLITE for tracerouting. The default port is 53.
- `-w waittime` — (Optional) Enter the time in seconds to wait for a response to a probe. The default is 5 seconds.
- `-z sendwait` — (Optional) Enter the minimal time interval to wait between probes. The default is 0. A value greater than 10 specifies a number in milliseconds, otherwise it specifies a number of seconds. This option is useful when routers rate-limit ICMP messages.
- `--mtu` — (Optional) Discovers the maximum transmission unit (MTU) from the path being traced.
- `--back` — (Optional) Prints the number of backward hops when different from the forward direction.
- `host` — (Required) Enter the name or IP address of the destination device.
- `packet_len` — (Optional) Enter the total size of the probing packet. The default is 60 bytes for IPv4 and 80 for IPv6.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# traceroute www.dell.com
traceroute to www.dell.com (23.73.112.54), 30 hops max, 60 byte packets
 1  10.11.97.254 (10.11.97.254)  4.298 ms  4.417 ms  4.398 ms
 2  10.11.3.254 (10.11.3.254)  2.121 ms  2.326 ms  2.550 ms
 3  10.11.27.254 (10.11.27.254)  2.233 ms  2.207 ms  2.391 ms
 4  Host65.hbms.com (63.80.56.65)  3.583 ms  3.776 ms  3.757 ms
 5  host33.30.198.65 (65.198.30.33)  3.758 ms  4.286 ms  4.221 ms
 6  3.GigabitEthernet3-3.GW3.SCL2.ALTER.NET (152.179.99.173)  4.428 ms  2.593
ms  3.243 ms
 7  0.xe-7-0-1.XL3.SJC7.ALTER.NET (152.63.48.254)  3.915 ms  3.603 ms  3.790 ms
 8  TenGigE0-4-0-5.GW6.SJC7.ALTER.NET (152.63.49.254)  11.781 ms  10.600 ms
9.402 ms
 9  23.73.112.54 (23.73.112.54)  3.606 ms  3.542 ms  3.773 ms
```

Example (IPv6)

```
OS10# traceroute 20::1
traceroute to 20::1 (20::1), 30 hops max, 80 byte packets
 1  20::1 (20::1)  2.622 ms  2.649 ms  2.964 ms
```

Supported Releases 10.2.0E or later

unlock

Unlocks a previously locked candidate configuration file.

Syntax `unlock`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# unlock
```

Supported Releases 10.2.0E or later

write

Copies the current running configuration to the startup configuration file.

Syntax	`write {memory}`
Parameters	`memory` — Copy the current running configuration to the startup configuration.
Default	Not configured
Command Mode	EXEC
Usage Information	This command has the same effect as the `copy running-configuration startup-configuration` command. The running configuration is not saved to a local configuration file other than the startup configuration. Use the `copy` command to save running configuration changes to a local file.
Example	`OS10# write memory`
Supported Releases	10.2.0E or later

2

Interfaces

You can configure and monitor physical interfaces (Ethernet), port-channels, and virtual local area networks (VLANs) in Layer 2 (L2) or Layer 3 (L3) modes.

Table 1. Interface types

Interface type	Supported	Default mode	Requires creation	Default status
Ethernet (PHY)	L2, L3	unset	No	no shutdown enabled
Management	N/A	N/A	No	no shutdown enabled
Loopback	L3	L3	Yes	no shutdown enabled
Port-channel	L2, L3	unset	Yes	no shutdown enabled
VLAN	L2, L3	L3	Yes, except default	no shutdown enabled

Ethernet interfaces

Ethernet port interfaces are enabled by default. To disable an Ethernet interface, use the `shutdown` command.

To re-enable a disabled interface, use the `no shutdown` command.

1 Configure an Ethernet port interface from Global CONFIGURATION mode.

```
interface ethernet node/slot/port[:subport]
```

2 Disable and re-enable the Ethernet port interface in INTERFACE mode.

```
shutdown
```

```
no shutdown
```

Disable Ethernet port interface

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# shutdown
```

Enable Ethernet port interface

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
```

Unified port groups

In an OS10 unified port group, all ports operate in either Ethernet or Fibre Channel (FC) mode. You cannot mix modes for ports in the same unified port group. To activate Ethernet interfaces, configure a port group to operate in Ethernet mode and specify the port speed. To activate Fibre Channel interfaces, see Fibre Channel interfaces.

S4148U-ON

On the S4148U-ON switch, the available Ethernet and Fibre Channel interfaces in a port group depend on the currently configured port profile. For more information, see S4148U-ON port profiles.

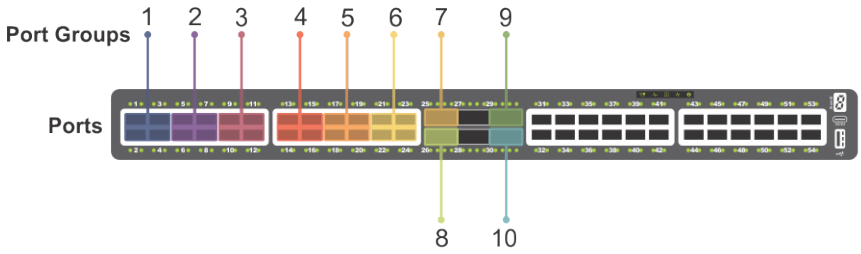


Figure 1. S4148U-ON unified port groups

To enable Ethernet interfaces in a unified port group:

1 Configure a unified port group in CONFIGURATION mode. Enter `1/1` for *node/slot*. The port-group range depends on the switch.

```
port-group node/slot/port-group
```

2 Activate the unified port group for Ethernet operation in PORT-GROUP mode. To activate a unified port group in Fibre Channel mode, see Fibre Channel interfaces. The available options depend on the switch.

```
mode Eth {100g-1x | 50g-2x | 40g-1x | 25g-4x | 10g-4x}
```

- `100g-1x` — Reset a port group to 100GE mode.
- `50g-2x` — Split a port group into two 50GE interfaces.
- `40g-1x` — Set a port group to 40GE mode for use with a QSFP+ 40GE transceiver.
- `25g-4x` — Split a port group into four 25GE interfaces.
- `10g-4x` — Split a port group into four 10GE interfaces.

3 Return to CONFIGURATION mode.

```
exit
```

4 Enter Ethernet Interface mode to configure other settings. Enter a single interface, a hyphen-separated range, or multiple interfaces separated by commas.

```
interface ethernet node/slot/port[:subport]
```

Configure Ethernet unified port interface

```
OS10(config)# port-group 1/1/13
OS10(conf-pg-1/1/13)# mode Eth 25g-4x
OS10(conf-pg-1/1/13)# exit
OS10(config)# interface ethernet 1/1/41:1
OS10(conf-if-eth1/1/41:1)#
```

View Ethernet unified port interface

```
OS10(config)# interface ethernet 1/1/41
OS10(conf-if-eth1/1/41:1)# show configuration
!
interface ethernet1/1/41:1
 no shutdown
```

L2 mode configuration

Each physical Ethernet interface uses a unique MAC address. Port-channels and VLANs use a single MAC address. By default, all the interfaces operate in L2 mode. From L2 mode you can configure switching and L2 protocols, such as VLANs and Spanning-Tree Protocol (STP) on an interface.

Enable L2 switching on a port interface in Access or Trunk mode. By default, an interface is configured in Access mode. Access mode allows L2 switching of untagged traffic on a single VLAN (VLAN 1 is the default). Trunk mode enables L2 switching of untagged traffic on the Access VLAN, and tagged traffic on one or more VLANs.

By default, native VLAN of a port is the default VLAN ID of the switch. You can change the native VLAN using the `switchport access vlan vlan-id` command.

A Trunk interface carries VLAN traffic that is tagged using 802.1q encapsulation. If an Access interface receives a packet with an 802.1q tag in the header that is different from the Access VLAN ID, it drops the packet.

By default, a trunk interface carries only untagged traffic on the Access VLAN. You must manually configure other VLANs for tagged traffic.

1 Select one of the two available options:
   - Configure L2 trunking in INTERFACE mode and the tagged VLAN traffic that the port can transmit. By default, a trunk port is not added to any tagged VLAN. You must create a VLAN before you can assign the interface to it.

     ```
     switchport mode trunk
     switchport trunk allowed vlan vlan-id-list
     ```
   - Reconfigure the access VLAN assigned to a L2 access or trunk port in INTERFACE mode.

     ```
     switchport access vlan vlan-id
     ```
2 Enable the interface for L2 traffic transmission in INTERFACE mode.

   ```
   no shutdown
   ```

L2 interface configuration

```
OS10(config)# interface ethernet 1/1/7
OS10(conf-if-eth1/1/7)# switchport mode trunk
OS10(conf-if-eth1/1/7)# switchport trunk allowed vlan 5,10
OS10(conf-if-eth1/1/7)# no shutdown
```

L3 mode configuration

Ethernet and port-channel interfaces are in L2 access mode by default. When you disablethe L2 mode and then assign an IP address to an Ethernet port interface, you place the port in L3 mode.

Configure one primary IP address in L3 mode. You can configure up to 255 secondary IP addresses on an interface. At least one interface in the system must be in L3 mode before you configure or enter a L3–protocol mode, such as OSPF.

1 Remove a port from L2 switching in INTERFACE mode.

   ```
   no switchport
   ```
2 Configure L3 routing in INTERFACE mode. Add `secondary` to configure backup IP addresses.

   ```
   ip address address [secondary]
   ```
3 Enable the interface for L3 traffic transmission in INTERFACE mode.

   ```
   no shutdown
   ```

L3 interface configuration

```
OS10(config)# interface ethernet 1/1/9
OS10(conf-if-eth1/1/9)# no switchport
OS10(conf-if-eth1/1/9)# ip address 10.10.1.92/24
OS10(conf-if-eth1/1/9)# no shutdown
```

View L3 configuration error

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ip address 1.1.1.1/24
% Error: remove Layer 2 configuration before assigning an IP
```

Fibre Channel interfaces

OS10 unified port groups support FC interfaces. A unified port group operates in Fibre Channel or Ethernet mode. To activate FC interfaces, configure a port group to operate in Fibre Channel mode and specify the port speed. By default, FC interfaces are disabled.

S4148U-ON

On a S4148U-ON switch, FC interfaces are available in all port groups. The activated FC interfaces depend on the currently configured port profile. For more information, see S4148U-ON port profiles.

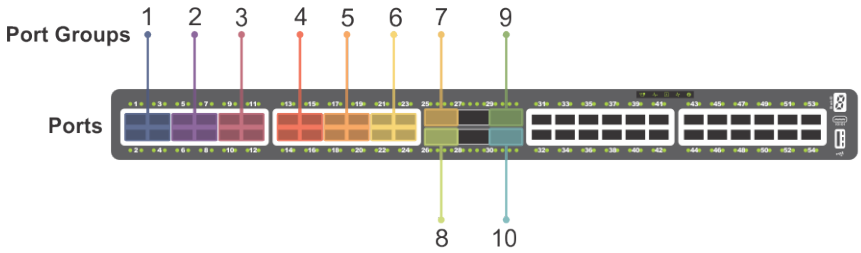


Figure 2. S4148U-ON unified port groups

1 Configure a unified port group in CONFIGURATION mode. Enter `1/1` for *node/slot*. The port-group range depends on the switch.

```
port-group node/slot/port-group
```

2 Activate the unified port group for FC operation in PORT-GROUP mode. The available FC modes depend on the switch.

```
mode fc {32g-4x | 32g-2x | 32g-1x | 16g-4x | 16g-2x | 8g-4x}
```

- `8g-4x` — Split a unified port group into four 8 GFC interfaces.
- `16g-2x` — Split a unified port group into two 16 GFC interfaces.
- `16g-4x` — Split a unified port group into four 16 GFC interfaces.
- `32g-1x` — Split a unified port group into one 32 GFC interface. A 1x-32G interface has a rate limit of 28G.
- `32g-2x` — Split a unified port group into two 32 GFC interfaces.
- `32g-4x` — Split a unified port group into four 32 GFC interfaces. Each 4x-32GE breakout interface has a rate limit of 25G.

3 Return to CONFIGURATION mode.

```
exit
```

4 Enter FC Interface mode to enable data transmission. Enter a single interface, a hyphen-separated range, or multiple interfaces separated by commas.

```
interface fibrechannel node/slot/port[:subport]
```

5 (Optional) Reconfigure the interface speed in INTERFACE mode.

```
speed {8 | 16 | 32 | auto}
```

6 Enable the FC interface in INTERFACE mode.

```
no shutdown
```

Configure FC interface

```
OS10(config)# port-group 1/1/15
OS10(conf-pg-1/1/15)# mode FC 16g-4x
OS10(conf-pg-1/1/15)# exit
OS10(config)# interface fibrechannel 1/1/43:1
OS10(conf-if-fc-1/1/43:1)# speed 32
OS10(conf-if-fc-1/1/43:1)# no shutdown
```

View FC interface

```
OS10(config)# interface fibrechannel 1/1/43:1
OS10(conf-if-fc-1/1/43:1)# show configuration
!
interface fibrechannel1/1/43:1
 no shutdown
 speed 32
 vfabric 100
```

```
OS10# show interface fibrechannel 1/1/43:1
Fibrechannel 1/1/43:1 is up, FC link is up
Address is 14:18:77:20:8d:fc, Current address is 14:18:77:20:8d:fc
```

```
Pluggable media present, QSFP+ type is QSFP+ 4x(16GBASE FC SW)
    Wavelength is 850
    Receive power reading is 0.0
FC MTU 2188 bytes
LineSpeed 8G
Port type is F, Max BB credit is 1
WWN is 20:78:14:18:77:20:8d:cf
Last clearing of "show interface" counters: 00:02:32
Input statistics:
     33 frames, 3508 bytes
     0 class 2 good frames, 33 class 3 good frames
     0 frame too long, 0 frame truncated, 0 CRC
     1 link fail, 0 sync loss
     0 primitive seq err, 0 LIP count
     0 BB credit 0, 0 BB credit 0 packet drops
Output statistics:
     33 frames, 2344 bytes
     0 class 2 frames, 33 class 3 frames
     0 BB credit 0, 0 oversize frames
6356027325 total errors
Rate Info:
     Input 116 bytes/sec, 1 frames/sec, 0% of line rate
     Output 78 bytes/sec, 1 frames/sec, 0% of line rate
Time since last interface status change: 00:00:24
```

Management interface

The Management interface provides OOB management access to the network device. You can configure the Management interface, but the configuration options on this interface are limited. You cannot configure gateway addresses and IP addresses if it appears in the main routing table. Proxy ARP is not supported on this interface.

1 Configure the Management interface in CONFIGURATION mode.

```
interface mgmt 1/1/1
```

2 By default, DHCP client is enabled on the Management interface. Disable the DHCP client operations in INTERFACE mode.

```
no ip address dhcp
```

3 Configure an IP address and mask on the Management interface in INTERFACE mode.

```
ip address A.B.C.D/prefix-length
```

4 Enable the Management interface in INTERFACE mode.

```
no shutdown
```

Configure management interface

```
OS10(config)# interface mgmt 1/1/1
OS10(conf-if-ma-1/1/1)# no ip address dhcp
OS10(conf-if-ma-1/1/1)# ip address 10.1.1.10/24
OS10(conf-if-ma-1/1/1)# no shutdown
```

VLAN interfaces

VLANs are logical interfaces and are, by default, in L2 mode. Physical interfaces and port-channels can be members of VLANs.

OS10 supports inter-VLAN routing. You can add IP addresses to VLANs and use them in routing protocols in the same manner that physical interfaces are used.

When using VLANs in a routing protocol, you must configure the `no shutdown` command to enable the VLAN for routing traffic. In VLANs, the `shutdown` command prevents L3 traffic from passing through the interface. L2 traffic is unaffected by this command.

- Configure an IP address in A.B.C.D/x format on the interface in INTERFACE mode. The secondary IP address is the interface's backup IP address.

```
ip address ip-address/mask [secondary]
```

Configure VLAN

```
OS10(config)# interface vlan 10
OS10(conf-if-vl-10)# ip address 1.1.1.2/24
```

You cannot simultaneously use egress rate shaping and ingress rate policing on the same VLAN.

User-configured default VLAN

By default, VLAN1 serves as the default VLAN for switching untagged L2 traffic on OS10 ports in Trunk or Access mode. The default VLAN is used for untagged protocol traffic sent and received between switches, such as STPs. If you use VLAN1 for data traffic for network-specific needs, reconfigure the VLAN ID of the default VLAN.

- Assign a new VLAN ID to the default VLAN in CONFIGURATION mode, from 1 to 4093.

  ```
  default vlan-id vlan-id
  ```

In the `show vlan` output, an asterisk (*) indicates the default VLAN.

Reconfigure default VLAN

```
OS10# show vlan
Q: A - Access (Untagged), T - Tagged
    NUM    Status    Description                    Q Ports
*   1      up                                       A Eth1/1/1-1/1/25,1/1/29,1/1/31-1/1/54

OS10(config)# interface vlan 10
Sep 19 17:28:10 OS10 dn_ifm[932]: Node.1-Unit.1:PRI:notice [os10:notify],
%Dell EMC (OS10) %IFM_ASTATE_UP: Interface admin state up :vlan10
OS10(conf-if-vl-10)# exit

OS10(config)# default vlan-id 10
Sep 19 17:28:15 OS10 dn_ifm[932]: Node.1-Unit.1:PRI:notice [os10:trap], %Dell EMC (OS10)
%IFM_OSTATE_DN: Interface operational state is down :vlan1
Sep 19 17:28:16 OS10 dn_ifm[932]: Node.1-Unit.1:PRI:notice [os10:trap], %Dell EMC (OS10)
%IFM_OSTATE_UP: Interface operational state is up :vlan10

OS10(config)# do show vlan
Codes: * - Default VLAN, M - Management VLAN, R - Remote Port Mirroring VLANs
Q: A - Access (Untagged), T - Tagged
    NUM    Status    Description                    Q Ports
    1      down
*   10     up                                       A Eth1/1/1-1/1/25,1/1/29,1/1/31-1/1/54
```

VLAN scale profile

When you scale the number of VLANs on a switch, use the VLAN scale profile. This consumes less memory. Enable the scale profile before you configure VLANs on the switch. The scale profile globally applies L2 mode on all VLANs you create and disables L3 transmission. To enable L3 routing traffic on a VLAN, use the `mode L3` command.

1. Configure the L2 VLAN scale profile in CONFIGURATION mode.

   ```
   scale-profile vlan
   ```

2. (Optional) Enable L3 routing on a VLAN in INTERFACE VLAN mode.

   ```
   mode L3
   ```

After you upgrade OS10 from an earlier version with configured VLANs, if you configure the VLAN scale profile and enable L3 routing on VLANs, save the configuration and use the `reload` command to apply the scale profile settings.

Apply VLAN scale profile

```
OS10(config)# scale-profile vlan
OS10(config)# interface vlan 10
OS10(conf-if-vl-10)# mode L3
```

Loopback interfaces

A Loopback interface is a virtual interface where the software emulates an interface. Because a Loopback interface is not associated to physical hardware entities, the Loopback interface status is not affected by hardware status changes.

Packets routed to a Loopback interface process locally to the OS10 device. Because this interface is not a physical interface, to provide protocol stability you can configure routing protocols on this interface. You can place Loopback interfaces in default L3 mode.

- Enter the Loopback interface number in CONFIGURATION mode, from 0 to 16383.
  ```
  interface loopback number
  ```
- Enter the Loopback interface number to view the configuration in EXEC mode.
  ```
  show interface loopback number
  ```
- Enter the Loopback interface number to delete a Loopback interface in CONFIGURATION mode.
  ```
  no interface loopback number
  ```

View Loopback interface

```
OS10# show interface loopback 4
Loopback 4 is up, line protocol is up
Hardware is unknown.
Interface index is 102863300
Internet address is 120.120.120.120/24
Mode of IPv4 Address Assignment : MANUAL
MTU 1532 bytes
Flowcontrol rx false tx false
ARP type: ARPA, ARP Timeout: 240
Last clearing of "show interface" counters : 00:00:11
Queuing strategy : fifo
    Input 0 packets, 0 bytes, 0 multicast
    Received 0 errors, 0 discarded
    Output 0 packets, 0 bytes, 0 multicast
    Output 0 errors, Output 0 invalid protocol
Time since last interface status change : 00:00:11
```

Port-channel interfaces

Port-channels are not configured by default. Link aggregation (LA) is a method of grouping multiple physical interfaces into a single logical interface — a link aggregation group (LAG) or port-channel. A port-channel aggregates the bandwidth of member links, provides redundancy, and load balances traffic. If a member port fails, the OS10 device redirects traffic to the remaining ports.

A physical interface can belong to only one port-channel at a time. A port-channel must contain interfaces of the same interface type and speed. OS10 supports a maximum of 128 port-channels, with up to thirty-two ports per channel.

To configure a port-channel, use the same configuration commands as the Ethernet port interfaces. Port-channels are transparent to network configurations and manage as a single interface. For example, configure one IP address for the group, and use the IP address for all routed traffic on the port-channel.

By configuring port channels, you can create larger capacity interfaces by aggregating a group of lower-speed links. For example, you can build a 40G interface by aggregating four 10G Ethernet interfaces together. If one of the four interfaces fails, traffic redistributes across the three remaining interfaces.

Static	Port-channels are statically configured.
Dynamic	Port-channels are dynamically configured using Link Aggregation Control Protocol (LACP).

Member ports of a LAG are added and programmed into the hardware based on the port ID, instead of the order the ports come up. Load balancing yields predictable results across resets and reloads.

Create port-channel

You can create a maximum of 128 port-channels, with up to 32 port members per group. Configure a port-channel similarly to a physical interface, enable or configure protocols, or ACLs to a port channel. After you enable the port-channel, place it in L2 or L3 mode.

To place the port-channel in L2 mode or configure an IP address to place the port-channel in L3 mode, use the `switchport` command.

- Create a port-channel in CONFIGURATION mode.
  ```
  interface port-channel id-number
  ```

Create port-channel

```
OS10(config)# interface port-channel 10
```

Add port member

When you add an interface to a port-channel:

- The administrative status applies to the port-channel.
- The port-channel configuration is applied to the member interfaces.
- A port-channel operates in either L2 (default) or L3 mode. To place a port-channel in L2 mode, use the `switchport mode` command. To place a port-channel in L3 mode and remove L2 configuration before you configure an IP address, use the `no switchport` command.
- All interfaces must have the same speed.
- An interface must not contain non-default L2/L3 configuration settings. Only the `description` and `shutdown` or `no shutdown`commands are supported. You cannot add an IP address or static MAC address to a member interface.
- You cannot enable flow control on a port-channel interface. Flow control is supported on physical interfaces that are port-channel members.
- Port-channels support 802.3ad LACP. LACP identifies similarly configured links and dynamically groups ports into a logical channel. LACP activates the maximum number of compatible ports that the switch supports in a port-channel.
- If you globally disable a spanning-tree operation, L2 interfaces that are LACP-enabled port-channel members may flap due to packet loops.

Add port member — static LAG

A static port-channel LAG contains member interfaces that you manually assign using the `channel-group mode on` command.

```
OS10(config)# interface port-channel 10
Aug 24 4:5:38: %Node.1-Unit.1:PRI:OS10 %dn_ifm
%log-notice:IFM_ASTATE_UP: Interface admin state up.:port-channel10
Aug 24 4:5:38: %Node.1-Unit.1:PRI:OS10 %dn_ifm
%log-notice:IFM_OSTATE_DN: Interface operational state is down.:port-channel10
OS10(conf-if-po-10)# exit

OS10(config)# interface ethernet 1/1/2
OS10(conf-if-eth1/1/2)# channel-group 10 mode on
Aug 24 4:5:56: %Node.1-Unit.1:PRI:OS10 %dn_ifm
%log-notice:IFM_OSTATE_UP: Interface operational state is up.:port-channel10
```

Add port member — dynamic LACP

LACP enables ports to dynamically bundle as members of a port-channel. To configure a port for LACP operation, use the `channel-group mode {active|passive}` command. Active and Passive modes allow LACP to negotiate between ports to determine if they can form a port channel based on their configuration settings.

```
OS10(config)# interface port-channel 100
OS10(conf-if-po-100)# exit
```

```
OS10(config)# interface ethernet 1/1/2
OS10(conf-if-eth1/1/2)# channel-group 100 mode active
```

Minimum links

Configure minimum links in a port-channel LAG that must be in *oper up* status to consider the port-channel to be in *oper up* status.

- Enter the number of links in a LAG that must be in *oper up* status in PORT-CHANNEL mode, from 1 to 32, default 1.

  ```
  minimum-links number
  ```

Configure minimum operationally up links

```
OS10(config)# interface port-channel 1
OS10(conf-if-po-1)# minimum-links 5
```

Assign Port Channel IP Address

You can assign an IP address to a port channel and use port channels in L3 routing protocols.

- Configure an IP address and mask on the interface in INTERFACE PORT-CHANNEL mode.

  ```
  ip address ip-address/mask [secondary-ip-address]
  ```

  - *ip-address/mask* — Specify an IP address in dotted-decimal A.B.C.D format and the mask.
  - *secondary-ip-address* — Specify a secondary IP address in dotted-decimal A.B.C.D format, which acts as the interface's backup IP address.

  Assign Port Channel IP Address

  ```
  OS10# configure terminal
  OS10(config)# interface port-channel 1
  OS10(conf-if-po-1)# ip address 1.1.1.1/24
  OS10(conf-if-po-1)#
  ```

Remove or disable port-channel

You can delete or disable a port-channel.

1. Delete a port-channel in CONFIGURATION mode.

   ```
   no interface port-channel channel-number
   ```

2. Disable a port-channel to place all interfaces within the port-channel operationally down in CONFIGURATION mode.

   ```
   shutdown
   ```

Delete port-channel

```
OS10(config)# interface port-channel 10
OS10(conf-if-po-10)# no interface port-channel 10
```

Load balance traffic

Use hashing to load balance traffic across member interfaces of a port-channel. Load balancing uses source and destination packet information to distribute traffic over multiple interfaces when transferring data to a destination.

For packets without an L3 header, OS10 automatically uses the `load-balancing mac—selection destination-mac` command for hash algorithms by default.

When you configure an IP and MAC hashing scheme at the same time, the MAC hashing scheme takes precedence over the IP hashing scheme.

- Select one or more methods of load balancing and replace the default IP 4-tuple method of balancing traffic over a port-channel in CONFIGURATION mode.

```
OS10(config)# load-balancing
  ingress-port       Ingress port configurations
  tcp-udp-selection  TCP-UDP port for load-balancing configurations
  ip-selection       IPV4 load-balancing configurations
  ipv6-selection     IPV6 load-balancing configurations
  mac-selection      MAC load-balancing configurations
```

  - `ingress-port [enable]` — Enables the ingress port configuration. This option is not supported on S5148F-ON.
  - `tcp-upd-selection [l4-destination-port | l4-source-port]` — Uses the Layer 4 destination IP address or Layer 4 source IP address in the hash calculation. This option is not supported on the S5148F-ON switch.
  - `ip-selection [destination-ip | source-ip | protocol | vlan-id | l4-destination-port | l4-source-port]` — Uses the destination IP address, source IP address, protocol, VLAN ID, Layer 4 destination IP address or Layer 4 source IP address in the hash calculation.
  - `ipv6-selection [destination-ip | source-ip | protocol | vlan-id | l4-destination-port | l4-source-port]` — Uses the destination IPv6 address, source IPv6 address, protocol, VLAN ID, Layer 4 destination IPv6 address or Layer 4 source IPv6 address in the hash calculation.
  - `mac—selection [destination-mac | source-mac] [ethertype | vlan-id]` — Uses the destination MAC address or source MAC address, and ethertype, or VLAN ID in the hash calculation.

Configure load balancing

```
OS10(config)# load-balancing ip-selection destination-ip source-ip
```

Change hash algorithm

The `load-balancing` command selects the hash criteria applied to traffic load balancing on port-channels. If you do not obtain even traffic distribution, use the `hash-algorithm` command to select the hash scheme for LAG. Rotate or shift the L2-bit LAG hash until you achieve the desired traffic distribution.

- Change the default (`0`) to another algorithm and apply it to LAG hashing in CONFIGURATION mode.

```
hash-algorithm lag crc
```

Change hash algorithm

```
OS10(config)# hash-algorithm lag crc
```

Configure interface ranges

Bulk interface configuration allows you to apply the same configuration to multiple physical or logical interfaces, or to display their current configuration. An interface range is a set of interfaces that you apply the same command to.

You can use interface ranges for:

- Ethernet physical interfaces
- Port channels
- VLAN interfaces

A bulk configuration includes any non-existing interfaces in an interface range from the configuration.

You can configure a default VLAN only if the interface range being configured consists of only VLAN ports. When a configuration in one of the VLAN ports fails, all the VLAN ports in the interface range are affected.

Create an interface range allowing other commands to be applied to that interface range using the `interface range` command.

Configure range of Ethernet addresses and enable them

```
OS10(config)# interface range ethernet 1/1/1-1/1/5
OS10(conf-range-eth1/1/1-1/1/5)# no shutdown
```

View the configuration

```
OS10(conf-range-eth1/1/1-1/1/5)# show configuration
!
interface ethernet1/1/1
 no shutdown
 switchport access vlan 1
!
interface ethernet1/1/2
 no shutdown
 switchport access vlan 1
!
interface ethernet1/1/3
 no shutdown
 switchport access vlan 1
!
interface ethernet1/1/4
 no shutdown
 switchport access vlan 1
!
interface ethernet1/1/5
 no shutdown
 switchport access vlan 1
```

Configure range of VLANs

```
OS10(config)# interface range vlan 1-100
OS10(conf-range-vl-1-100)#
```

Configure range of port channels

```
OS10(config)# interface range port-channel 1-25
OS10(conf-range-po-1-25)#
```

Switch-port profiles

A port profile determines the enabled front-panel ports and supported breakout modes on Ethernet and unified ports. Change the port profile on a switch to customize uplink and unified port operation, and the availability of front-panel data ports.

To change the port profile at the next reboot, use the `switch-port-profile` command with the desired profile, save it to the startup configuration, and use the `reload` command to apply the changes.

1 Configure a platform-specific port profile in CONFIGURATION mode. For a standalone switch, enter `1/1` for `node/unit`.

```
switch-port-profile node/unit profile
```

2 Save the port profile change to the startup configuration in EXEC mode.

```
write memory
```

3 Reload the switch in EXEC mode.

```
reload
```

The switch reboots with the new port configuration and resets the system defaults, except for the switch-port profile and these configured settings:

- Management interface 1/1/1 configuration

- Management IPv4/IPv6 static routes
- System hostname
- Unified Forwarding Table (UFT) mode
- ECMP maximum paths

You must manually reconfigure other settings on a switch after you apply a new port profile and reload the switch.

NOTE: After you change the switch-port profile, do not immediately back up and restore the startup file without using the `write memory` command and reloading the switch using the `reload` command. Otherwise, the new profile does not take effect.

Configure port profile

```
OS10(config)# switch-port-profile 1/1 profile-6
OS10(config)# exit
OS10# write memory
OS10# reload
```

Verify port profile

```
OS10(config)# show switch-port-profile 1/1

|  Node/Unit  |  Current    |  Next-boot  | Default   |
|-------------+-------------------+-------------------|
|     1/1     |  profile-2  |  profile-2  | profile-1 |

Supported Profiles:
profile-1
profile-2
profile-3
profile-4
profile-5
profile-6
```

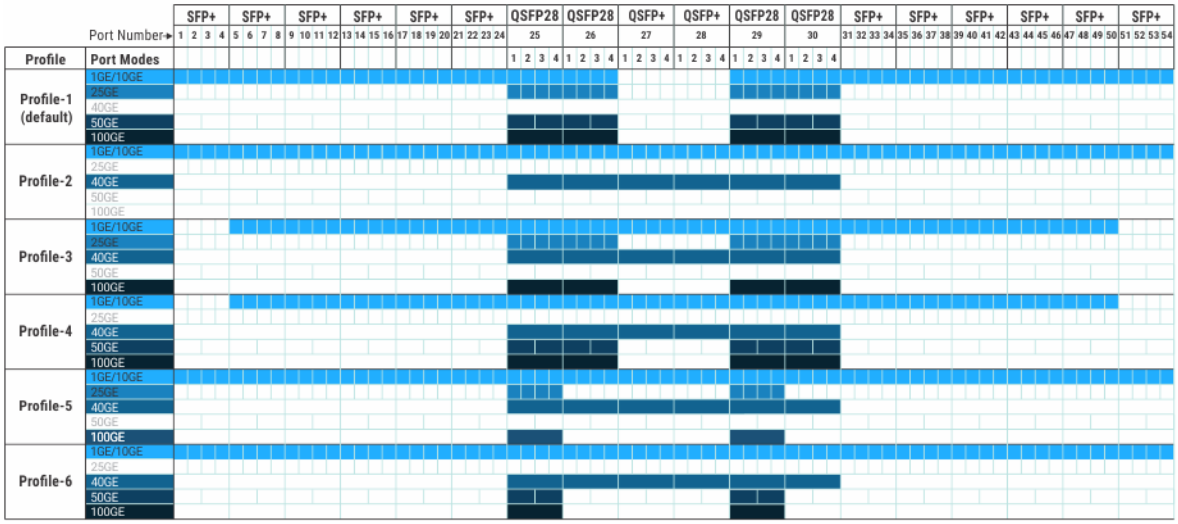
S4148-ON Series port profiles

On the S4148-ON Series of switches, port profiles determine the available front-panel Ethernet ports and supported breakout interfaces on uplink ports. In the port profile illustration, blue boxes indicate the supported ports and breakout interfaces. Blank spaces indicate ports and speeds that are not available.

- 10GE mode is an SFP+ 10GE port or a 4x10G breakout of a QSFP+ or QSFP28 port.
- 25GE is a 4x25G breakout of a QSFP28 port.
- 40GE mode is a QSFP+ port or a QSFP28 port that supports QSFP+ 40GE transceivers.
- 50GE is a 2x50G breakout of a QSFP28 port.
- 100GE mode is a QSFP28 port.

NOTE: For S4148U-ON port profiles with both unified and Ethernet ports, see S4148U-ON port profiles. An S4148U-ON unified port supports Fibre Channel and Ethernet modes.

For example, `profile-1` enables 10G speed on forty-eight ports (1-24 and 31-54), and 4x10G breakouts on QSFP28 ports 25-26 and 29-30; QSFP+ ports 27 and 28 are deactivated. `profile-3` enables 10G speed on forty ports, and 4x10G breakouts on all QSFP28 and QSFP+ ports. Similarly, `profile-1` disables 40G speed on ports 25-30; `profile-3` enables 40G on these ports. For more information, see switch-port-profile.



1GE mode: 1GE is supported only on SFP+ ports; 1GE is not supported on QSFP+ and QSFP28 ports 25-26.

Breakout interfaces: Use the `interface breakout` command in Configuration mode to configure 4x10G, 4x25G, and 2x50G breakout interfaces.

To view the ports that belong to each port group, use the `show port-group` command.

S4148U-ON port profiles

S4148U-ON port profiles determine the available front-panel unified and Ethernet ports and supported breakout interfaces. In the port profile illustration, blue boxes indicate the supported Ethernet port modes and breakout interfaces. Brown boxes indicate the supported Fibre Channel port modes and breakout interfaces. Blank spaces indicate ports and speeds that are not available. Unified port groups are numbered 1 to 10.

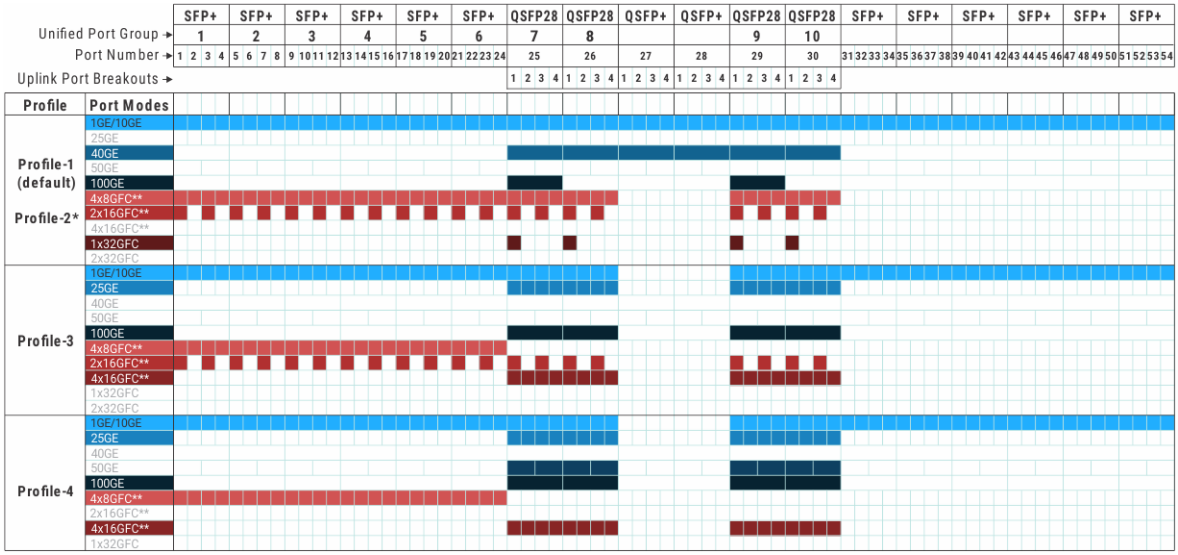
S4148U-ON unified port modes—SFP+ ports 1-24 and QSFP28 ports 25-26 and 29-30:

- 10GE is an SFP+ port in Ethernet mode or a 4x10G breakout of a QSFP+ or QSFP28 port in Ethernet mode.
- 25GE is a 4x25G breakout of a QSFP28 Ethernet port.
- 40GE is a QSFP+ or QSFP28 Ethernet port that uses QSFP+ 40GE transceivers.
- 50GE is a 2x50G breakout of a QSFP28 Ethernet port.
- 100GE is a QSFP28 Ethernet port.
- 4x8GFC are breakout interfaces in an SFP+ or QSFP28 FC port group.
- 2x16GFC are breakout interfaces (subports 1 and 3) in an SFP+ or QSFP28 FC port group.
- 4x16GFC are breakout interfaces in a QSFP28 FC port group.
- 1x32GFC (subport 1) are breakout interfaces in a QSFP28 FC port group.

S4148U-ON Ethernet modes—QSFP+ ports 27-28 and SFP+ ports 31-54:

- 10GE mode is an SFP+ 10GE port or a 4x10G breakout of a QSFP+ port.
- 40GE mode is a QSFP+ port.

For example, all S4148U-ON activate support 10G speed on unified ports 1-24 and Ethernet ports 31-54, but only `profile-1` and `profile-2` activate QSFP+ ports 27-28 in 40GE mode with 4x10G breakouts. Similarly, all S4148U-ON profiles activate 8GFC speed on unified ports 1-24, but only `profile-1`, `profile-2`, and `profile-3` activate 2x16GFC in port groups 1-6. In QSFP28 port groups, `profile-1` and `profile-2` support 1x32GFC; `profile-3` and `profile-4` support 4x16GFC.



*`profile-1` and `profile-2` activate the same port mode capability on unified and Ethernet ports. The difference is that in `profile-1`, by default SFP+ unified ports 1-24 come up in Fibre Channel mode with 2x16GFC breakouts per port group. In `profile-2`, by default SFP+ unified ports 1-24 come up in Ethernet 10GE mode. `profile-1` allows you to connect FC devices for plug-and-play; `profile-2` is designed for a standard Ethernet-based data network.

****Oversubscription**: Configure oversubscription to support bursty storage traffic on a Fibre Channel interface. Oversubscription allows a port to operate faster, but may result in traffic loss. To support oversubscription, use the `speed` command in Interface Configuration mode. This command is not supported on an Ethernet interface. In S4148U-ON port profiles:

- SFP+ and QSFP28 port groups in 4x8GFC mode support 16GFC oversubscription on member interfaces.
- QSFP28 ports in 2x16GFC mode support 32GFC oversubscription. SFP+ port groups in 2x16GFC mode do not support 32GFC oversubscription. 2x16GFC mode activates subports 1 and 3.
- QSFP28 ports in 4x16GFC mode support 32GFC oversubscription.

Breakout interfaces:

- To configure breakout interfaces on a unified port, use the `mode {FC | Eth}` command in Port-Group Configuration mode. The `mode {FC | Eth}` command configures a unified port to operate at line rate and guarantees no traffic loss.
- To configure breakout interfaces on a QSFP+ Ethernet port, use the `interface breakout` command in global Configuration mode.

1GE mode: Only SFP+ ports support 1GE; QSFP+ and QSFP28 ports 25 to 30 do not support 1GE.

To view the ports that belong to each port group, use the `show port-group` command.

Configure breakout mode

Using a supported breakout cable, you can split a 40GE QSFP+ or 100GE QSFP28 Ethernet port into separate breakout interfaces. All breakout interfaces have the same speed. You can set a QSFP28 port to operate in 40GE mode with a QSFP+ transceiver.

```
interface breakout node/slot/port map {10g-4x | 25g-4x | 40g-1x | 50g-2x |
100g-1x}
```

- *`node/slot/port`* — Enter the physical port information.
- `10g-4x` — Split a QSFP28 or QSFP+ port into four 10G interfaces.
- `25g-4x` — Split a QSFP28 port into four 25G interfaces.
- `40g-1x` — Set a QSFP28 port to use with a QSFP+ 40G transceiver.
- `50g-2x` — Split a QSFP28 port into two 50G interfaces.

- `100g-1x` — Reset a QSFP28 port to 100G speed.

To configure an Ethernet breakout interface, use the `interface ethernet` *node/slot/port:subport* command in CONFIGURATION mode.

Each breakout interface operates at the configured speed. Use the `no` version of the `interface breakout` command to reset a port to its default speed: 40G or 100G.

To configure breakout interfaces on a unified port, use the `mode {Eth | FC}` command in Port-Group Configuration mode.

Configure interface breakout

```
OS10(config)# interface breakout 1/1/7 map 10g-4x
```

Display interface breakout

```
OS10# show interface status

--------------------------------------------------------------------
Port       Description  Status Speed Duplex Mode Vlan Tagged-Vlans
--------------------------------------------------------------------
Eth 1/1/1               down   0     auto    -
Eth 1/1/2               down   0     auto    A    1    -
Eth 1/1/7:1             down   0     auto    A    1    -
Eth 1/1/7:2             down   0     auto    A    1    -
Eth 1/1/7:3             down   0     auto    A    1    -
Eth 1/1/7:4             down   0     auto    A    1    -
Eth 1/1/25              down   0     auto    A    1    -
```

Breakout auto-configuration

You can globally enable front-panel Ethernet ports to automatically detect SFP pluggable media in a QSFP+ or QSFP28 port. The port autoconfigures breakout interfaces for media type and speed. For example, if you plug a 40G direct attach cable (DAC) with 4x10G far-side transceivers into a QSFP28 port, the port autoconfigures in `10g-4x` Interface-breakout mode.

RJ-45 ports and ports that are members of a port group do not support breakout auto-configuration. Breakout auto-configuration is disabled by default.

Enable breakout auto-configuration

```
OS10(config)# feature auto-breakout
```

Display breakout auto-configuration

Before you plug a cable in Ethernet port 1/1/25:

```
OS10# show interface status

--------------------------------------------------------------------
Port       Description  Status Speed Duplex Mode Vlan Tagged-Vlans
--------------------------------------------------------------------
Eth 1/1/1               down   0     auto    -
Eth 1/1/2               down   0     auto    A    1    -
Eth 1/1/25              down   0     auto    A    1    -
Eth 1/1/29              down   0     auto    A    1    -
```

After you enter `feature auto-breakout` and plug a breakout cable in Ethernet port 1/1/25:

```
OS10# show interface status

--------------------------------------------------------------------
Port       Description  Status Speed Duplex Mode Vlan Tagged-Vlans
--------------------------------------------------------------------
Eth 1/1/1               down   0     auto    -
```

```
Eth 1/1/2                down    0      auto     A     1     -
Eth 1/1/25:1             down    0      auto     A     1     -
Eth 1/1/25:2             down    0      auto     A     1     -
Eth 1/1/25:3             down    0      auto     A     1     -
Eth 1/1/25:4             down    0      auto     A     1     -
Eth 1/1/29               down    0      auto     A     1     -
```

Forward error correction

Forward error correction (FEC) enhances data reliability.

FEC modes supported in OS10:

- CL74-FC — Supports 25G
- CL91-RS — Supports 100G
- CL108-RS — Supports 25G
- off — Disables FEC

NOTE: OS10 does not support FEC on 10G and 40G.

Configure FEC

```
OS10(config)# interface ethernet 1/1/41
OS10(conf-if-eth1/1/41)# fec CL91-RS
```

View FEC configuration

```
OS10# show interface ethernet 1/1/41
Ethernet 1/1/41 is up, line protocol is up
Hardware is Dell EMC Eth, address is e4:f0:04:3e:1a:06
    Current address is e4:f0:04:3e:1a:06
Pluggable media present, QSFP28 type is QSFP28_100GBASE_CR4_2M
    Wavelength is 64
    Receive power reading is
Interface index is 17306108
Internet address is not set
Mode of IPv4 Address Assignment: not set
Interface IPv6 oper status: Disabled
MTU 1532 bytes, IP MTU 1500 bytes
LineSpeed 100G, Auto-Negotiation on
FEC is cl91-rs, Current FEC is cl91-rs
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 60
Last clearing of "show interface" counters: 00:00:17
Queuing strategy: fifo
Input statistics:
     7 packets, 818 octets
     2 64-byte pkts, 0 over 64-byte pkts, 5 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     7 Multicasts, 0 Broadcasts, 0 Unicasts
     0 runts, 0 giants, 0 throttles
     0 CRC, 0 overrun, 0 discarded
Output statistics:
     15 packets, 1330 octets
     10 64-byte pkts, 0 over 64-byte pkts, 5 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     15 Multicasts, 0 Broadcasts, 0 Unicasts
     0 throttles, 0 discarded, 0 Collisions, 0 wred drops
Rate Info(interval 30 seconds):
     Input 0 Mbits/sec, 0 packets/sec, 0% of line rate
     Output 0 Mbits/sec, 0 packets/sec, 0% of line rate
Time since last interface status change: 00:00:13
--more--
```

Energy-efficient Ethernet

Energy-efficient Ethernet (EEE) reduces power consumption of physical layer devices (PHYs) during idle periods. EEE allows Dell EMC Networking devices to conform to green computing standards.

An Ethernet link consumes power when a link is idle. EEE allows Ethernet links to use Regular Power mode only during data transmission. EEE is enabled on devices that support LOW POWER IDLE (LPI) mode. Such devices save power by entering LPI mode during periods when no data is transmitting.

In LPI mode, systems on both ends of the link saves power by shutting down certain services. EEE transitions into and out of LPI mode transparently to upper-layer protocols and applications.

EEE advertises during the auto-negotiation stage. Auto-negotiation detects abilities supported by the device at the other end of the link, determines common abilities, and configures joint operation.

Auto-negotiation performs at power-up, on command from the LAN controller, on detection of a PHY error, or following Ethernet cable re-connection. During the link establishment process, both link partners indicate their EEE capabilities. If EEE is supported by both link partners for the negotiated PHY type, EEE functions independently in either direction.

Changing the EEE configuration resets the interface because the device restarts Layer 1 auto-negotiation. You may want to enable Link Layer Discovery Protocol (LLDP) for devices that require longer wake-up times before they are able to accept data on their receive paths. Doing so enables the device to negotiate extended system wake-up times from the transmitting link partner.

Enable energy-efficient Ethernet

EEE is disabled by default. To reduce power consumption, enable EEE.

1 Enter the physical Ethernet interface information in CONFIGURATION mode.

```
interface ethernet node/slot/port[:subport]
```

2 Enable EEE in INTERFACE mode.

```
eee
```

Enable EEE

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# eee
```

Disable EEE

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no eee
```

Clear EEE counters

You can clear EEE counters on physical Ethernet interfaces globally or per interface.

Clear all EEE counters

```
OS10# clear counters interface eee
Clear all eee counters [confirm yes/no]:yes
```

Clear counters for specific interface

```
OS10# clear counters interface 1/1/48 eee
Clear eee counters on ethernet1/1/48 [confirm yes/no]:yes
```

View EEE status/statistics

You can view the EEE status or statistics for a specified interface, or all interfaces, using the `show` commands.

View EEE status for a specified interface

```
OS10# show interface ethernet 1/1/48 eee

Port            EEE  Status   Speed    Duplex
-----------------------------------------------
Eth 1/1/48      on   up       1000M
```

View EEE status on all interfaces

```
OS10# show interface eee

Port            EEE  Status   Speed    Duplex
-----------------------------------------------
Eth 1/1/1       off  up       1000M
...
Eth 1/1/47      on   up       1000M
Eth 1/1/48      on   up       1000M
Eth 1/1/49      n/a
Eth 1/1/50      n/a
Eth 1/1/51      n/a
Eth 1/1/52      n/a
```

View EEE statistics for a specified interface

```
OS10# show interface ethernet 1/1/48 eee statistics
Eth 1/1/48
  EEE             : on
  TxIdleTime(us)  : 2560
  TxWakeTime(us)  : 5
  Last Clearing   : 18:45:53
  TxEventCount    : 0
  TxDuration(us)  : 0
  RxEventCount    : 0
  RxDuration(us)  : 0
```

View EEE statistics on all interfaces

```
OS10# show interface eee statistics

Port        EEE  TxEventCount   TxDuration(us)  RxEventCount   RxDuration(us)
-------------------------------------------------------------------------------
Eth 1/1/1   off  0              0               0              0
...
Eth 1/1/47  on   0              0               0              0
Eth 1/1/48  on   0              0               0              0
Eth 1/1/49  n/a
...
Eth 1/1/52  n/a
```

EEE commands

clear counters interface eee

Clears all EEE counters.

Syntax	`clear counters interface eee`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	None
Example	`OS10# clear counters interface eee` `Clear all eee counters [confirm yes/no]:yes`
Supported Releases	10.3.0E or later

clear counters interface ethernet eee

Clears EEE counters on a specified Ethernet interface.

Syntax	`clear counters interface ethernet` *`node/slot/port[:subport]`* `eee`
Parameters	`node/slot/port[:subport]` —Enter the interface information.
Default	Not configured
Command Mode	EXEC
Usage Information	None
Example	`OS10# clear counters interface 1/1/48 eee` `Clear eee counters on ethernet1/1/48 [confirm yes/no]:yes`
Supported Releases	10.3.0E or later

eee

Enables or disables energy-efficient Ethernet (EEE) on physical ports.

Syntax	`eee`
Parameters	None
Default	Enabled on Base-T devices and disabled on S3048-ON and S4048T-ON switches.
Command Mode	Interface
Usage Information	To disable EEE, use the `no` version of this command.
Example (Enable EEE)	`OS10(config)# interface ethernet 1/1/1` `OS10(conf-if-eth1/1/1)# eee`
Example (Disable EEE)	`OS10(config)# interface ethernet 1/1/2` `OS10(conf-if-eth1/1/2)# no eee`
Supported Releases	10.3.0E or later

show interface eee

Displays the EEE status for all interfaces.

Syntax `show interface eee`

Parameters None

Default Not configured

Command Mode EXEC

Example

```
OS10# show interface eee

Port            EEE  Status   Speed    Duplex
-----------------------------------------------
Eth 1/1/1       off  up       1000M
...
Eth 1/1/47      on   up       1000M
Eth 1/1/48      on   up       1000M
Eth 1/1/49      n/a
Eth 1/1/50      n/a
Eth 1/1/51      n/a
Eth 1/1/52      n/a
```

Supported Releases 10.3.0E or later

show interface eee statistics

Displays EEE statistics for all interfaces.

Syntax `show interface eee statistics`

Parameters None

Default Not configured

Command Mode EXEC

Example

```
OS10# show interface eee statistics

Port        EEE  TxEventCount   TxDuration(us)  RxEventCount   RxDuration(us)
-------------------------------------------------------------------------------
Eth 1/1/1   off  0              0               0              0
...
Eth 1/1/47  on   0              0               0              0
Eth 1/1/48  on   0              0               0              0
Eth 1/1/49  n/a
...
Eth 1/1/52  n/a
```

Supported Releases 10.3.0E or later

show interface ethernet eee

Displays the EEE status for a specified interface.

Syntax `show interface ethernet` *node/slot/port[:subport]* `eee`

Parameters *node/slot/port[:subport]*—Enter the interface information.

Default	Not configured
Command Mode	EXEC

Example	`OS10# show interface ethernet 1/1/48 eee`
	
	`Port  EEE  Status  Speed  Duplex`
	`----------`
	`Eth 1/1/48  on  up  1000M`

Supported Releases	10.3.0E or later

show interface ethernet eee statistics

Displays EEE statistics for a specified interface.

Syntax	`show interface ethernet` *node/slot/port[:subport]* `eee statistics`
Parameters	*node/slot/port[:subport]*—Enter the interface information.
Default	Not configured
Command Mode	EXEC

Example	`OS10# show interface ethernet 1/1/48 eee statistics`
	`Eth 1/1/48`
	`  EEE : on`
	`  TxIdleTime(us) : 2560`
	`  TxWakeTime(us) : 5`
	`  Last Clearing : 18:45:53`
	`  TxEventCount : 0`
	`  TxDuration(us) : 0`
	`  RxEventCount : 0`
	`  RxDuration(us) : 0`

Supported Releases	10.3.0E or later

View interface configuration

To view basic interface information, use the `show interface`, `show running-configuration`, and `show interface status` commands. Stop scrolling output from a `show` command by entering CTRL+C. Display information about a physical or virtual interface in EXEC mode, including up/down status, MAC and IP addresses, and input/output traffic counters.

```
show interface [type]
```

- `phy-eth` *node/slot/port[:subport]* — Display information about physical media connected to the interface.
- `status` — Display interface status.
- `ethernet` *node/slot/port[:subport]* — Display Ethernet interface information.
- `loopback` *id* — Display Loopback interface information, from 0 to 16383.
- `mgmt` *node/slot/port* — Display Management interface information.
- `port-channel` *id-number* — Display port-channel interface information, from 1 to 128.
- `vlan` *vlan-id* — Display the VLAN interface information, from 1 to 4093.

View interface information

```
OS10# show interface
Ethernet 1/1/1 is up, line protocol is down
Hardware is Eth, address is 00:0c:29:66:6b:90
    Current address is 00:0c:29:66:6b:90
```

```
Pluggable media present, QSFP+ type is QSFP+ 40GBASE CR4
    Wavelength is 64
    Receive power reading is 0.000000 dBm
Interface index is 15
Internet address is not set
Mode of IPv4 Address Assignment: not set
Interface IPv6 oper status: Enabled
MTU 1532 bytes, IP MTU 1500 bytes
LineSpeed 0, Auto-Negotiation on
Configured FEC is off, Negotiated FEC is off
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 60
Last clearing of "show interface" counters: 02:46:35
Queuing strategy: fifo
Input statistics:
     0 packets, 0 octets
     0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     0 Multicasts, 0 Broadcasts, 0 Unicasts
     0 runts, 0 giants, 0 throttles
     0 CRC, 0 overrun, 0 discarded
Output statistics:
     0 packets, 0 octets
     0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     0 Multicasts, 0 Broadcasts, 0 Unicasts
     0 throttles, 0 discarded, 0 Collisions, 0 wred drops
Rate Info(interval 30 seconds):
     Input 0 Mbits/sec, 0 packets/sec, 0% of line rate
     Output 0 Mbits/sec, 0 packets/sec, 0% of line rate
Time since last interface status change: 02:46:36


Ethernet 1/1/2 is up, line protocol is up
Hardware is Eth, address is 00:0c:29:66:6b:94
    Current address is 00:0c:29:66:6b:94
Pluggable media present, QSFP+ type is QSFP+ 40GBASE CR4
    Wavelength is 64
    Receive power reading is 0.000000 dBm
Interface index is 17
Internet address is not set
Mode of IPv4 Address Assignment: not set
Interface IPv6 oper status: Enabled
Link local IPv6 address: fe80::20c:29ff:fe66:6b94/64
MTU 1532 bytes, IP MTU 1500 bytes
LineSpeed 40G, Auto-Negotiation on
Configured FEC is off, Negotiated FEC is off
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 60
Last clearing of "show interface" counters: 02:46:35
Queuing strategy: fifo
Input statistics:
     0 packets, 0 octets
     0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     0 Multicasts, 0 Broadcasts, 0 Unicasts
     0 runts, 0 giants, 0 throttles
     0 CRC, 0 overrun, 0 discarded
Output statistics:
     0 packets, 0 octets
     0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     0 Multicasts, 0 Broadcasts, 0 Unicasts
     0 throttles, 0 discarded, 0 Collisions, 0 wred drops
Rate Info(interval 30 seconds):
     Input 0 Mbits/sec, 0 packets/sec, 0% of line rate
     Output 0 Mbits/sec, 0 packets/sec, 0% of line rate
```

```
Time since last interface status change: 02:46:35
--more--
```

View specific interface information

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# show configuration
!
interface ethernet1/1/1
 ip address 1.1.1.1/24
 no switchport
 no shutdown
```

View candidate configuration

```
OS10(conf-if-eth1/1/1)# show configuration candidate
!
interface ethernet1/1/1
 ip address 1.1.1.1/24
 no switchport
 no shutdown
```

View running configuration

```
OS10# show running-configuration
Current Configuration ...
!
interface ethernet1/1/1
  no ip address
  shutdown
!
interface ethernet1/1/2
  no ip address
  shutdown
!
interface ethernet1/1/3
  no ip address
  shutdown
!
interface ethernet1/1/4
  no ip address
  shutdown
...
```

View L3 interfaces

```
OS10# show ip interface brief
Interface Name            IP-Address          OK       Method       Status     Protocol
================================================================================================
Ethernet 1/1/1            unassigned          NO       unset        up          down
Ethernet 1/1/2            unassigned          YES      unset        up          up
Ethernet 1/1/3            3.1.1.1/24          YES      manual       up          up
Ethernet 1/1/4            4.1.1.1/24          YES      manual       up          up
Ethernet 1/1/5            unassigned          NO       unset        up          down
Ethernet 1/1/6            unassigned          NO       unset        up          down
Ethernet 1/1/7            unassigned          NO       unset        up          down
Ethernet 1/1/8            unassigned          NO       unset        up          down
Ethernet 1/1/9            unassigned          NO       unset        up          down
Ethernet 1/1/10           unassigned          NO       unset        up          down
Ethernet 1/1/11           unassigned          NO       unset        up          down
Ethernet 1/1/12           unassigned          NO       unset        up          down
Ethernet 1/1/13           unassigned          NO       unset        up          down
Ethernet 1/1/14           unassigned          NO       unset        up          down
Ethernet 1/1/15           unassigned          NO       unset        up          down
Ethernet 1/1/16           unassigned          NO       unset        up          down
Ethernet 1/1/17           unassigned          NO       unset        up          down
Ethernet 1/1/18           unassigned          NO       unset        up          down
Ethernet 1/1/19           unassigned          NO       unset        up          down
Ethernet 1/1/20           unassigned          NO       unset        up          down
```

```
Ethernet 1/1/21           unassigned          NO       unset        up          down
Ethernet 1/1/22           unassigned          NO       unset        up          down
Ethernet 1/1/23           unassigned          NO       unset        up          down
Ethernet 1/1/24           unassigned          NO       unset        up          down
Ethernet 1/1/25           unassigned          NO       unset        up          down
Ethernet 1/1/26           unassigned          NO       unset        up          down
Ethernet 1/1/27           unassigned          NO       unset        up          down
Ethernet 1/1/28           unassigned          NO       unset        up          down
Ethernet 1/1/29           unassigned          NO       unset        up          down
Ethernet 1/1/30           unassigned          NO       unset        up          down
Ethernet 1/1/31           unassigned          NO       unset        up          down
Ethernet 1/1/32           unassigned          NO       unset        up          down
Management 1/1/1          10.16.153.226/24    YES      manual       up          up
Vlan 1                    unassigned          NO       unset        up          down
Vlan 10                   unassigned          NO       unset        up          down
Vlan 20                   unassigned          NO       unset        up          down
Vlan 30                   unassigned          NO       unset        up          down
```

View VLAN configuration

```
OS10# show vlan
Codes: * - Default VLAN, M - Management VLAN, R - Remote Port Mirroring VLANs,
       @ - Attached to Virtual Network
Q: A - Access (Untagged), T - Tagged
    NUM    Status    Description                    Q Ports
    1      Inactive                                 A Eth1/1/1,1/1/6-1/1/32
    10     Inactive
    20     Inactive
    22     Inactive
    23     Active                                   A Eth1/1/2
    24     Inactive
    25     Inactive
    26     Inactive
    27     Inactive
    28     Inactive
    29     Inactive
    30     Inactive
```

Digital optical monitoring

The digital optical monitoring (DOM) feature monitors the digital optical media for temperature, voltage, bias, transmission power (Tx), and reception power (Rx). This feature also generates event logs, alarms, and traps for any fluctuations, when configured thresholds are reached.

There are four threshold levels for each of the DOM categories—temperature, voltage, bias, transmission power, and reception power as summarized in the following table:

- High
- High warning
- Low
- Low warning

The OS10 DOM subsystem periodically monitors the optical transceivers for temperature, voltage, bias, transmission power and reception power changes and generate event logs, alarms, and traps when their respective values cross the predefined thresholds.

Table 2. DOM Alarms

Alarm Category	Alarm Name	Traps Generated?	Severity Level
Temperature	Temperature high	Y	Major
	Temperature high warning	N	Minor
	Temperature low	Y	Major

Alarm Category	Alarm Name	Traps Generated?	Severity Level
	Temperature low warning	N	Minor
Voltage	Voltage high	Y	Major
	Voltage high warning	N	Minor
	Voltage low	Y	Major
	Voltage low warning	N	Minor
Bias	Bias high	Y	Major
	Bias high warning	N	Minor
	Bias low	Y	Major
	Bias low warning	N	Minor
Power transmission (Tx)	Tx high	Y	Major
	Tx high warning	N	Minor
	Tx low	Y	Major
	Tx low warning	N	Minor
Power reception (Rx)	Rx high	Y	Major
	Rx high warning	N	Minor
	Rx low	Y	Major
	Rx low warning	N	Minor

You can enable or disable the DOM feature, configure traps, and view the DOM status.

Enable DOM and DOM traps

To generate DOM alarms, do the following.

1 Enable DOM.

```
OS10(config)# dom enable
```

2 Enable DOM traps.

```
OS10(config)# snmp-server enable traps dom
```

You can run the `show alarms` command in EXEC mode to view any alarms that are generated.

View DOM alarms

```
OS10# show alarms

Index   Severity  Name                  Raise-time              Source
-----   --------  --------------------  ----------------------- ------
0       major     EQM_MEDIA_TEMP_HIGH   Tue 06-04-2019 12:32:07  Node.1-Unit.1
```

View DOM event log message

The following are examples of event logs:

- High temperature warning:

```
Aug 03 06:35:47 OS10 dn_eqm[9135]: [os10:alarm], %Dell EMC (OS10)
%EQM_MEDIA_TEMP_HIGH: Media high temperature threshold crossed major warning
SET media 1/1/21 high threshold crossed, 82.00:78.00 Aug 03 06:35:47 OS10
dn_eqm[9135]: [os10:alarm], %Dell EMC (OS10) %EQM_MEDIA_VOLTAGE_HIGH: Media high
```

```
voltage threshold crossed major warning SET media 1/1/21 high threshold crossed,
6.00:3.63
```

In this example, the threshold for high temperature is 78.00, but the current temperature is 82.00.

- High reception power warning:

```
Aug 03 06:35:47 OS10 dn_eqm[9135]: [os10:alarm], %Dell EMC (OS10)
%EQM_MEDIA_RX_POWER_HIGH: Media high rx_power threshold crossed major warning
SET media 1/1/21 high threshold crossed, 7.00:3.30 Aug 03 06:35:47 OS10
dn_eqm[9135]: [os10:alarm], %Dell EMC (OS10) %EQM_MEDIA_BIAS_HIGH: Media high
bias threshold crossed major warning SET media 1/1/21 high threshold crossed,
120.00:105.00
```

In this example, the threshold for high reception power is 3.30, but the current reception power is 7.00.

View DOM traps

The following are examples of DOM traps.

```
2018-08-21 17:38:18 <UNKNOWN> [UDP: [10.11.56.49]:51635->[10.11.86.108]:162]:
iso.3.6.1.2.1.1.3.0 = Timeticks: (0) 0:00:00.00 iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.
3.6.1.4.1.674.11000.5000.100.4.1.3.1.15      iso.3.6.1.4.1.674.11000.5000.100.4.1.3.2.4 = INTEGER: 1 iso.
3.6.1.4.1.674.11000.5000.100.4.1.3.2.5 = INTEGER: 21        iso.3.6.1.4.1.674.11000.5000.100.4.1.3.2.1 = INTEGER:
1081393   iso.3.6.1.4.1.674.11000.5000.100.4.1.3.2.3 = INTEGER: 1 iso.3.6.1.4.1.674.11000.5000.100.4.1.3.2.2 = STRING:
"SET media 1/1/21 high threshold crossed, 82.00:78.00"
```

```
2018-08-21 17:38:18 <UNKNOWN> [UDP: [10.11.56.49]:48521->[10.11.86.108]:162]:
iso.3.6.1.2.1.1.3.0 = Timeticks: (1) 0:00:00.01 iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.
3.6.1.4.1.674.11000.5000.100.4.1.3.1.19      iso.3.6.1.4.1.674.11000.5000.100.4.1.3.2.3 =
INTEGER: 1 iso.3.6.1.4.1.674.11000.5000.100.4.1.3.2.1 = INTEGER: 1081397   iso.
3.6.1.4.1.674.11000.5000.100.4.1.3.2.4 = INTEGER: 1 iso.3.6.1.4.1.674.11000.5000.100.4.1.3.2.2
= STRING: "SET media 1/1/21 high threshold crossed, 6.00:3.63"    iso.
3.6.1.4.1.674.11000.5000.100.4.1.3.2.5 = INTEGER: 21
```

Interface commands

channel-group

Assigns an interface to a port-channel group.

Syntax `channel-group channel-number mode {active | on | passive}`

Parameters

- *channel-number* — Enter a port-channel number, from 1 to 128.
- `mode` — Sets LACP Actor mode.
- `active` — Sets Channeling mode to Active.
- `on` — Sets Channeling mode to static.
- `passive` — Sets Channeling mode to passive.

Default Not configured

Command Mode INTERFACE

Usage Information The `no` version of this command resets the value to the default, and unassigns the interface from the port-channel group.

Example

```
OS10(config)# interface ethernet 1/1/2:1
OS10(conf-if-eth1/1/2:1)# channel-group 20 mode active
```

Supported Releases 10.3.0E or later

default vlan-id

Reconfigures the VLAN ID of the default VLAN.

Syntax `default vlan-id vlan-id`

Parameters *vlan-id* — Enter the default VLAN ID number, from 1 to 4093.

Default VLAN1

Command Mode CONFIGURATION

Usage Information By default, VLAN1 serves as the default VLAN for switching untagged L2 traffic on OS10 ports in Trunk or Access mode. If you use VLAN1 for network-specific data traffic, reconfigure the VLAN ID of the default VLAN. The command reconfigures the access VLAN ID, the default VLAN, of all ports in Switchport Acess mode. Ensure that the VLAN ID exists before configuring it as the default VLAN.

Example

```
OS10(config)# default vlan-id 10

OS10(config)# do show running-configuration
...
!
interface vlan1
no shutdown
!
interface vlan10
no shutdown
!
interface ethernet1/1/1
no shutdown
switchport access vlan 10
!
interface ethernet1/1/2
no shutdown
switchport access vlan 10
!
interface ethernet1/1/3
no shutdown
switchport access vlan 10
!
interface ethernet1/1/4
no shutdown
switchport access vlan 10
```

Supported Releases 10.4.0E(R1) or later

description (Interface)

Configures a textual description of an interface.

Syntax `description string`

Parameters *string* — Enter a text string for the interface description. A maximum of 240 characters.

Default Not configured

Command Mode INTERFACE

Usage Information

- To use special characters as a part of the description string, enclose the string in double quotes.

- Spaces between characters are not preserved after entering this command unless you enclose the entire description in quotation marks; for example, "*text description*".
- Enter a text string after the `description` command to overwrite any previously configured text string.
- Use the `show running-configuration interface` command to view descriptions configured for each interface.
- The `no` version of this command deletes the description.

Example

```
OS10(conf-if-eth1/1/7)# description eth1/1/7
```

Supported Releases 10.2.0E or later

duplex

Configures Duplex mode on the Management port.

Syntax `duplex {full | half | auto}`

Parameters

- `full` — Set the physical interface to transmit in both directions.
- `half` — Set the physical interface to transmit in only one direction.
- `auto` — Set the port to auto-negotiate speed with a connected device.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information You can only use this command on the Management port. The `no` version of this command removes the duplex mode configuration from the management port.

Example

```
OS10(conf-if-ma-1/1/1)# duplex auto
```

Supported Releases 10.3.0E or later

enable dom

Enables or disables the DOM feature.

Syntax `dom enable`

Parameters None

Default Disabled

Command Mode CONFIGURATION

Usage Information The `no` version of this command disables digital optical monitoring.

Example

```
OS10# configure terminal
OS10(config)# dom enable
```

```
OS10# configure terminal
OS10(config)# no dom enable
```

Supported Releases 10.4.3.0 or later

enable dom traps

Enables DOM traps if the specified parameter crosses the defined threshold three times.

Syntax `snmp-server enable traps dom {temperature | voltage | rx-power | tx-power | bias}`

Parameters `temperature | voltage | rx-power | tx-power | bias` — Enter the keyword to enable DOM traps for the specified category.

Default Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command disables the DOM traps.

Example

```
OS10# configure terminal
OS10(config)# snmp-server enable traps dom temperature
```

```
OS10# configure terminal
OS10(config)# no snmp-server enable traps dom temperature
```

Supported Releases 10.4.3.0 or later

feature auto-breakout

Enables front-panel Ethernet ports to automatically detect SFP media and autoconfigure breakout interfaces.

Syntax `feature auto-breakout`

Parameters None

Default Not configured

Command mode CONFIGURATION

Usage information After you enter the `feature auto-breakout` command and plug a supported breakout cable in a QSFP+ or QSFP28 port, the port autoconfigures breakout interfaces for media type and speed.

Use the `interface breakout` command to manually configure breakout interfaces. The media type plugged into a port is no longer automatically learned. The `no` version of this command disables the auto-breakout feature.

Example

```
OS10(config)# feature auto-breakout
```

Supported releases 10.4.0E(R1) or later

fec

Configures Forward Error Correction on 25G and 100G interfaces.

Syntax `fec {CL74-FC | CL91-RS | CL108-RS | off}`

Parameters

- `CL74-FC` — Supports 25G
- `CL91-RS` — Supports 100G
- `CL108-RS` — Supports 25G

- `off` — Disables FEC

Defaults
- For 25G interfaces: `off`
- For 100G interfaces: `CL91-RS`

Command Mode CONFIGURATION

Usage Information The `no` version of this command resets the value to the default.

Example
```
OS10(config)# interface ethernet 1/1/41
OS10(conf-if-eth1/1/41)# fec CL91-RS
```

Supported Releases 10.3.0E or later

interface breakout

Splits a front-panel Ethernet port into multiple breakout interfaces.

Syntax `interface breakout` *`node/slot/port`* `map {100g-1x | 40g-1x | 25g-4x | 10g-4x | 25g-4x}`

Parameters
- *`node/slot/port`* — Enter the physical port information.
- `100g-1x` — Reset a QSFP28 port to 100G speed.
- `40g-1x` — Set a QSFP28 port to use with a QSFP+ 40GE transceiver.
- `25g-4x` — Split a QSFP28 port into four 25GE interfaces.
- `10g-4x` — Split a QSFP28 or QSFP+ port into four 10GE interfaces

Default Not configured

Command Mode CONFIGURATION

Usage Information
- Each breakout interface operates at the configured speed; for example, 10G or 25G.
- The `no interface breakout` *`node/slot/port`* command resets a port to its default speed: 40G or 100G.
- To configure breakout interfaces on a unified port, use the `mode {Eth | FC}` command in Port-Group Configuration mode.

Example
```
OS10(config)# interface breakout 1/1/41 map 10g-4x
```

Supported Releases 10.2.2E or later

interface ethernet

Configures a physical Ethernet interface.

Syntax `interface ethernet` *`node/slot/port:subport`*

Parameters *`node/slot/port:subport`* — Enter the Ethernet interface information.

Default Not configured

Command Mode	CONFIGURATION
Usage Information	The `no` version of this command deletes the interface.
Example	`OS10(config)# interface ethernet 1/1/10:1` `OS10(conf-if-eth1/1/10:1)#`
Supported Releases	10.2.0E or later

interface loopback

Configures a Loopback interface.

Syntax	`interface loopback` *id*
Parameters	*id* — Enter the Loopback interface ID number, from 0 to 16383.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command deletes the Loopback interface.
Example	`OS10(config)# interface loopback 100` `OS10(conf-if-lo-100)#`
Supported Releases	10.2.0E or later

interface mgmt

Configures the Management port.

Syntax	`interface mgmt` *node/slot/port*
Parameters	*node/slot/port* — Enter the physical port interface information for the Management interface.
Default	Enabled
Command Mode	CONFIGURATION
Usage Information	You cannot delete a Management port. To assign an IP address to the Management port, use the `ip address` command.
Example	`OS10(config)# interface mgmt 1/1/1` `OS10(conf-if-ma-1/1/1)#`
Supported Releases	10.2.0E or later

interface null

Configures a null interface on the switch.

Syntax	`interface null` *number*
Parameters	*number* — Enter the interface number to set as null (0).
Default	0

Command Mode	CONFIGURATION
Usage Information	You cannot delete the Null interface. The only configuration command possible in a Null interface is `ip unreachables`.
Example	`OS10(config)# interface null 0` `OS10(conf-if-nu-0)#`
Supported Releases	10.3.0E or later

interface port-channel

Creates a port-channel interface.

Syntax	`interface port-channel` *`channel-id`*
Parameters	*`channel-id`* — Enter the port-channel ID number, from 1 to 128.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command deletes the interface.
Example	`OS10(config)# interface port-channel 10` `OS10(conf-if-po-10)#`
Supported Releases	10.2.0E or later

interface range

Configures a range of Ethernet, port-channel, or VLAN interfaces for bulk configuration.

Syntax

`interface range {ethernet` *`node/slot/port[:subport]-node/slot/port[:subport]`*`, [...]} | {port-channel` *`IDnumber-IDnumber`*`,[ ...]} | vlan` *`vlanID-vlanID`*`,[...]}`

Parameters

- *`node/slot/port[:subport]-node/slot/port[:subport]`* — Enter a range of Ethernet interfaces.
- *`IDnumber-IDnumber`* — Enter a range of port-channel numbers, from 1 to 128.
- *`vlanID-vlanID`* — Enter a range VLAN ID numbers, from 1 to 4093.

Default	Not configured
Command Mode	CONFIGURATION

Usage Information

Enter up to six comma-separated interface ranges without spaces between commas. When creating an interface range, interfaces are not sorted and appear in the order entered. You cannot mix interface configuration such as Ethernet ports with VLANs.

- Abulk configuration is created if at least one interface is valid.
- Non-existing interfaces are excluded from the bulk configuration with a warning message.
- This command has multiple port ranges, the prompt excludes the smaller port range.
- If you enter overlapping port ranges, the port range extends to the smallest port and the largest end port.
- You can only use VLAN and port-channel interfaces created using the `interface vlan` and `interface port-channel` commands.

- You cannot create virtual VLAN or port-channel interfaces using the `interface range` command.
- The `no` version of this command deletes the interface range.

Example

```
OS10(config)# interface range ethernet 1/1/7-1/1/24
OS10(conf-range-eth1/1/7-1/1/24)#
```

Supported Releases 10.2.0E or later

interface vlan

Creates a VLAN interface.

Syntax `interface vlan vlan-id`

Parameters *vlan-id* — Enter the VLAN ID number, from 1 to 4093.

Default VLAN 1

Command Mode CONFIGURATION

Usage Information FTP, TFTP, MAC ACLs, and SNMP operations are not supported. IP ACLs are supported on VLANs only. The `no` version of this command deletes the interface.

Example

```
OS10(config)# interface vlan 10
OS10(conf-if-vl-10)#
```

Supported Releases 10.2.0E or later

link-bundle-utilization

Configures link-bundle utilization.

Syntax `link-bundle-utilization trigger-threshold value`

Parameters *value* — Enter the percentage of port-channel bandwidth that triggers traffic monitoring on port-channel members, from 0 to 100.

Default Disabled

Command Mode CONFIGURATION

Usage Information None

Example

```
OS10(config)# link-bundle-utilization trigger-threshold 10
```

Supported Releases 10.2.0E or later

mode

Configures a front-panel unified port group to operate in Fibre Channel or Ethernet mode, or a QSFP28-DD or QSFP28 port group to operate in Ethernet mode, with the specified speed on activated interfaces.

Syntax

```
mode {Eth {100g-2x | 100g-1x | 50g-2x | 40g-2x | 40g-1x | 25g-8x [fabric-
expander-mode] | 25g-4x | 10g-8x | 10g-4x} | FC {32g-4x | 32g-2x | 32g-1x |
16g-4x | 16g-2x | 8g-4x}}
```

Parameters

- `mode Eth` — Configure a port group in Ethernet mode and set the speed to:
  - `100g-2x` — Split a QSFP28-DD port into two 100GE interfaces.
  - `100g-1x` — Reset a QSFP28 port group to 100GE mode.
  - `50g-2x` — Split a QSFP28 port into two 50GE interfaces.
  - `40g-2x` — Split a port group into two 40GE interfaces.
  - `40g-1x` — Set a port group to 40G mode for use with a QSFP+ 40GE transceiver.
  - `25g-8x fabric-expander-mode` — Split a QSFP28-DD port into eight 25GE interfaces for connection to a Fabric Expander.
  - `25g-8x` — Split a port group into eight 25GE interfaces.
  - `25g-4x` — Split a port group into four 25GE interfaces.
  - `10g-8x` — Split a port group into eight 10GE interfaces.
  - `10g-4x` — Split a port group into four 10GE interfaces.
- `mode FC` — Configure a port group in Fibre Channel mode and set the speed to:
  - `32g-4x` — Split a port group into four 32GFC interfaces.
  - `32g-2x` — Split a port group into two 32GFC interfaces, subports 1 and 3.
  - `32g-1x` — Split a port group into one 32GFC interface, subport 1.
  - `16g-4x` — Split a port group into four 16GFC interfaces; supports 4x32GFC oversubscription.
  - `16g-2x` — Split a port group into two 16GFC interfaces using ports 1 and 3.
  - `8g-4x` — Split a port group into four 8GFC interfaces.

Default

S4148U-ON: Depends on the port profile activated.

MX9116n Fabric Switching Engine:

- QSFP28-DD port groups 1 to 9 operate in 8x25GE fabric-expander mode (FEM).
- QSFP28-DD port groups 10 to 12 operate in 2x100GE mode.
- QSFP28 port groups 13 and 14 operate in 1x100GE mode.
- Unified port groups 15 and 16 operate in ethernet 1x100GE mode.

Command Mode

PORT-GROUP

Usage Information

- The `mode {FC | Eth}` command configures a port group to operate at line rate and guarantees no traffic loss.
- To configure oversubscription on a FC interface, use the `speed` command.
- To configure breakout interfaces on an Ethernet port, use the `interface breakout` command.
- To view the currently active ports and subports, use the `show interfaces status` command.
- The `no` version of the command resets port-group interfaces to the default Ethernet port mode/speed. Use the `no mode` command before you reset the mode on an interface.

Example

```
OS10(conf-pg-1/1/2)# mode FC 16g-4x

OS10(conf-pg-1/1/8)# mode Eth 10g-4x
```

Example: Reset mode

```
OS10(conf-pg-1/1/2)# mode FC 16g-4x
OS10(conf-pg-1/1/2)# no mode
OS10(conf-pg-1/1/2)# mode Eth 10g-4x
```

Supported Releases

10.3.1E or later

mode l3

Enables L3 routing on a VLAN after you configure the VLAN scale profile.

Syntax `mode l3`

Parameters None

Defaults Not configured

Command Mode INTERFACE VLAN

Usage Information To configure the VLAN scale profile, use the `scale-profile vlan` command. The scale profile globally applies L2 mode on all VLANs you create and disables L3 transmission. To enable L3 routing traffic on a VLAN, use the `mode L3` command.

Example

```
OS10(config)# interface vlan 10
OS10(conf-if-vl-10)# mode L3
```

Supported Releases 10.4.0E(X2) or later

mtu

Sets the link maximum transmission unit (MTU) frame size for an Ethernet L2 or L3 interface.

Syntax `mtu value`

Parameters *value* — Enter the maximum frame size in bytes, from 1280 to 65535.

Default 1532 bytes

Command Mode INTERFACE

Usage Information To return to the default MTU value, use the `no mtu` command. If an IP packet includes a L2 header, the IP MTU must be at least 32 bytes smaller than the L2 MTU.

- Port-channels
  - All members must have the same link MTU value and the same IP MTU value.
  - The port channel link MTU and IP MTU must be less than or equal to the link MTU and IP MTU values you configure on the channel members. For example, if the members have a link MTU of 2100 and an IP MTU 2000, the port channel's MTU values cannot be higher than 2100 for link MTU or 2000 bytes for IP MTU.
- VLANS
  - All members of a VLAN must have same IP MTU value.
  - Members can have different link MTU values. Tagged members must have a link MTU four bytes higher than untagged members to account for the packet tag.
  - The VLAN link MTU and IP MTU must be less than or equal to the link MTU and IP MTU value you configure on the VLAN members. For example, the VLAN contains tagged members with a link MTU of 1522 and IP MTU of 1500 and untagged members with link MTU of 1518 and IP MTU of 1500. The VLAN's link MTU cannot be higher than 1518 bytes and its IP MTU cannot be higher than 1500 bytes.

Example

```
OS10(conf-if-eth1/1/7)# mtu 3000
```

Supported Releases 10.2.0E or later

port-group

Configures a group of front-panel unified ports, or a double-density QSFP28 (QSFP28-DD) or single-density QSFP28 port group.

Syntax `port-group node/slot/port-group`

Parameters

- *node/slot* — Enter `1/1` for *node/slot* when you configure a port group.
- *port-group* — Enter the port-group number, from 1 to 16. The available port-group range depends on the switch.

Default Not configured

Command mode CONFIGURATION

Usage information Enter PORT-GROUP mode to:

- Configure unified ports in Fibre Channel or Ethernet mode and break out interfaces with a specified speed.
- Break out an MX9116n QSFP28-DD or QSFP28 port group into multiple interfaces with a specified speed.

To view the ports that belong to a port group, use the `show port-group` command.

Example

```
OS10(config)# port-group 1/1/8
OS10(conf-pg-1/1/8)#
```

Supported releases 10.3.1E or later

scale-profile vlan

Configures the L2 VLAN scale profile on a switch.

Syntax `scale-profile vlan`

Parameters None

Defaults Not configured

Command Mode CONFIGURATION

Usage Information Use the VLAN scale profile when you scale the number of VLANs so that the switch consumes less memory. Enable the scale profile before you configure VLANs on the switch. The scale profile globally applies L2 mode on all VLANs you create and disables L3 transmission. The `no` version of the command disables L2 VLAN scaling. To enable L3 routing traffic on a VLAN, use the `mode L3` command.

Example

```
OS10(config)# scale-profile vlan
```

Supported Releases 10.4.0E(X2) or later

show interface

Displays interface information.

Syntax `show interface [type]`

Parameters `interface type` — Enter the interface type:

- `phy-eth node/slot/port[:subport]` — Display information about physical ports connected to the interface.
- `status` — Display interface status.
- `ethernet node/slot/port[:subport]` — Display Ethernet interface information.
- `loopback id` — Display Loopback IDs, from 0 to 16383.
- `mgmt node/slot/port` — Display Management interface information.
- `null` — Display null interface information.
- `port-channel id-number` — Display port channel interface IDs, from 1 to 128.
- `vlan vlan-id` — Display the VLAN interface number, from 1 to 4093.

Default Not configured

Command Mode EXEC

Usage Information Use the `do show interface` command to view interface information from other command modes.

Example

```
OS10# show interface
Ethernet 1/1/2 is up, line protocol is up
Hardware is Dell EMC Eth, address is 00:0c:29:54:c8:57
    Current address is 00:0c:29:54:c8:57
Pluggable media present, QSFP+ type is QSFP+ 40GBASE CR 1.0M
    Wavelength is 64
    Receive power reading is 0.0
Interface index is 17305094
Internet address is not set
Mode of IPv4 Address Assignment: not set
Interface IPv6 oper status: Enabled
Link local IPv6 address: fe80::20c:29ff:fe54:c857/64
Global IPv6 address: 2::1/64
MTU 1532 bytes, IP MTU 1500 bytes
LineSpeed 40G, Auto-Negotiation on
FEC is auto, Current FEC is off
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 60
Last clearing of "show interface" counters: 00:40:14
Queuing strategy: fifo
Input statistics:
     0 packets, 0 octets
     0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     0 Multicasts, 0 Broadcasts, 0 Unicasts
     0 runts, 0 giants, 0 throttles
     0 CRC, 0 overrun, 0 discarded
Output statistics:
     0 packets, 0 octets
     0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     0 Multicasts, 0 Broadcasts, 0 Unicasts
     0 throttles, 0 discarded, 0 Collisions, 0 wreddrops
Rate Info(interval 299 seconds):
     Input 0 Mbits/sec, 0 packets/sec, 0% of line rate
     Output 0 Mbits/sec, 0 packets/sec, 0% of line rate
Time since last interface status change: 3 weeks 1 day 20:30:38

--more--
```

Example (port channel)

```
OS10# show interface port-channel 1
Port-channel 1 is up, line protocol is down
Address is 90:b1:1c:f4:a5:8c, Current address is 90:b1:1c:f4:a5:8c
Interface index is 85886081
Internet address is not set
Mode of IPv4 Address Assignment: not set
MTU 1532 bytes
LineSpeed 0
```

```
Minimum number of links to bring Port-channel up is 1
Maximum active members that are allowed in the portchannel is 5
Members in this channel:
ARP type: ARPA, ARP Timeout: 60
```

```
OS10# show interface port-channel summary
LAG Mode Status Uptime Ports
22 L2 up 20:38:08 Eth 1/1/10 (Up)
                  Eth 1/1/11 (Down)
                  Eth 1/1/12 (Inact)
23 L2 up 20:34:32 Eth 1/1/20 (Up)
                  Eth 1/1/21 (Up)
                  Eth 1/1/22 (Up)
```

Supported Releases 10.2.0E or later

show inventory media

Displays installed media in switch ports.

Syntax `show inventory media`

Parameters None

Command Mode EXEC

Usage Information Use the `show inventory media` command to verify the media type inserted in a port.

Example

```
OS10# show inventory media
-----------------------------------------------------------------------
               System Inventory Media
-----------------------------------------------------------------------
Node/Slot/Port  Category     Media                Serial    Dell EMC
                                                  Number    Qualified
-----------------------------------------------------------------------
1/1/1           Not Present
1/1/2           SFP+         SFP+ 10GBASE SR      AM70843   true
1/1/3           Not Present
1/1/4           SFP+         SFP+ 10GBASE SR      AKN0LC7   false
1/1/5           SFP+         SFP+ 10GBASE SR      AM718GQ   true
1/1/6           SFP+         SFP+ 10GBASE SR      AM708XM   true
1/1/7           SFP+         SFP+ 10GBASE SR      AQ2237K   true
1/1/8           SFP+         SFP+ 10GBASE SR      AGT047N   true
1/1/9           Not Present
1/1/10          Not Present
1/1/11          Not Present
1/1/12          Not Present
1/1/13          Not Present
1/1/14          Not Present
1/1/15          SFP+         SFP+ 10GBASE SR      AK60QJN   false
1/1/16          SFP+         SFP+ 10GBASE SR      AL30KWM   true
1/1/17          SFP+         SFP+ 10GBASE SR      AQ22DMB   true
1/1/18          SFP+         SFP+ 10GBASE SR      AQM146U   true
...
```

Supported Releases 10.2.0E or later

show link-bundle-utilization

Displays information about the link-bundle utilization.

Syntax `show link-bundle-utilization`

Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	None

Example

```
OS10# show link-bundle-utilization

Link-bundle trigger threshold - 60
```

Supported Releases	10.2.0E or later

show port-channel summary

Displays port-channel summary information.

Syntax	`show port-channel summary`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	None

Example

```
OS10(conf-if-eth1/1/4)# do show port-channel summary
Flags: D - Down I - member up but inactive P - member up and active
U - Up (port-channel)

Group Port-Channel Type Protocol Member Ports

22 port-channel22 (U) Eth STATIC 1/1/2(D) 1/1/3(P)
23 port-channel23 (D) Eth DYNAMIC 1/1/4(I)
```

Example (Interface)

```
OS10(conf-range-eth1/1/10-1/1/11,1/1/13,1/1/14)# do show port-channel summary
Flags: D - Down U - member up but inactive P - member up and active
U - Up (port-channel)

Group Port-Channel Type Protocol Member Ports

22 port-channel22 (U) Eth STATIC 1/1/10(P) 1/1/11(P) 1/1/12(P) 1/1/13(P)
1/1/14(P) 1/1/15(P) 1/1/16(P) 1/1/17(P) 1/1/18(P) 1/1/19(P)
23 port-channel23 (D) Eth STATIC
OS10(config)# interface range e1/1/12-1/1/13,1/1/15,1/1/17-1/1/18
OS10(conf-range-eth1/1/12-1/1/13,1/1/15,1/1/17-1/1/18)# no channel-group
OS10(conf-range-eth1/1/12-1/1/13,1/1/15,1/1/17-1/1/18)# do show port-channel
summary
Flags: D - Down U - member up but inactive P - member up and active
U - Up (port-channel)

Group Port-Channel Type Protocol Member Ports

22 port-channel22 (U) Eth STATIC 1/1/10(P) 1/1/11(P) 1/1/14(P) 1/1/16(P)
1/1/19(P)
23 port-channel23 (D) Eth STATIC
```

Supported Releases	10.2.0E or later

show port-group

Displays the current port-group configuration on a switch.

Syntax `show port-group`

Parameters None

Default None

Command Mode EXEC

Usage Information To view the ports that belong to each port group, use the `show port-group` command. To configure a port group, use the `port-group` command.

Example: S4148U-ON

```
OS10(config)# show port-group
port-group mode          ports
1/1/1      Eth 10g-4x    1 2 3 4
1/1/2      FC  16g-2x    5 6 7 8
1/1/3      FC  16g-2x    9 10 11 12
1/1/4      FC  16g-2x    13 14 15 16
1/1/5      FC  16g-2x    17 18 19 20
1/1/6      FC  16g-2x    21 22 23 24
1/1/7      Eth 100g-1x   25
1/1/8      Eth 40g-1x    26
1/1/9      Eth 100g-1x   29
1/1/10     Eth 40g-1x    30
```

Example: MX9116n Fabric Engine

```
OS10(config)# show port-group
Port-group          Mode          Ports     FEM
port-group1/1/1     Eth 25g-8x    17  18    FEM
port-group1/1/2     Eth 25g-8x    19  20    FEM
port-group1/1/3     Eth 25g-8x    21  22    FEM
port-group1/1/4     Eth 25g-8x    23  24    FEM
port-group1/1/5     Eth 25g-8x    25  26    FEM
port-group1/1/6     Eth 25g-8x    27  28    FEM
port-group1/1/7     Eth 25g-8x    29  30    FEM
port-group1/1/8     Eth 25g-8x    31  32    FEM
port-group1/1/9     Eth 25g-8x    33  34    FEM
port-group1/1/10    Eth 100g-2x   35  36     -
port-group1/1/11    Eth 100g-2x   37  38     -
port-group1/1/12    Eth 100g-2x   39  40     -
port-group1/1/13    Eth 100g-1x   41         -
port-group1/1/14    Eth 100g-1x   42         -
port-group1/1/15    Eth 100g-1x   43         -
port-group1/1/16    Eth 100g-1x   44         -
```

Supported Releases 10.3.1E or later

show switch-port-profile

Displays the current and default port profile on a switch.

Syntax `show switch-port-profile` *node/slot*

Parameters

- *node/slot* — Enter the switch information. For a standalone switch, enter `1/1`.

Default `profile-1`

Command Mode EXEC

Usage Information A switch-port profile determines the available front-panel ports and breakout modes on Ethernet and unified ports. To display the current port profile, use the `show switch-port-profile` command. To reset the switch to the default port profile, use the `no switch-port-profile` *node/slot* command.

Example

```
OS10(config)# show switch-port-profile 1/1

|  Node/Unit  |  Current    |  Next-boot  | Default   |
|-------------+-------------------+-------------------|
|     1/1     |  profile-2  |  profile-2  | profile-1 |

Supported Profiles:
profile-1
profile-2
profile-3
profile-4
profile-5
profile-6
```

Supported Releases 10.3.1E or later

show system

Displays the status of the DOM feature, whether it is enabled or disabled.

Syntax `show system`

Parameters None

Defaults DOM disabled

Command Mode EXEC

Usage Information None

Example

```
OS10# show system
Node Id : 1
MAC : 90:b1:1c:f4:aa:b2
Number of MACs : 129
Up Time : 02:08:43

-- Unit 1 --
Status : up
System Identifier : 1
Down Reason : user-triggered
Digital Optical Monitoring : disable
```

Supported Releases 10.4.3.0 or later

show vlan

Displays the current VLAN configuration.

Syntax `show vlan [`*vlan-id*`]`

Parameters *vlan-id* — (Optional) Enter a VLAN ID, from 1 to 4093.

Default Not configured

Command Mode EXEC

Usage Information	None

Example	`OS10# show vlan`
	`Codes: * - Default VLAN, M - Management VLAN, R - Remote Port Mirroring VLANs`
	`Q: A - Access (Untagged), T - Tagged`
	`NUM Status Description Q Ports`
	`1 down`

Supported Releases	10.2.0E or later

shutdown

Disables an interface.

Syntax	`shutdown`
Parameters	None
Default	Disabled
Command Mode	INTERFACE
Usage Information	This command marks a physical interface as unavailable for traffic. Disabling a VLAN or a port-channel causes different behavior. When you disable a VLAN, the L3 functions within that VLAN are disabled, and L2 traffic continues to flow. Use the `shutdown` command on a port-channel to disable all traffic on the port-channel, and the individual interfaces. Use the `no shutdown` command to enable a port-channel on the interface. The `shutdown` and `description` commands are the only commands that you can configure on an interface that is a port-channel member.
Example	`OS10(config)# interface ethernet 1/1/7` `OS10(conf-if-eth1/1/7)# no shutdown`
Supported Releases	10.2.0E or later

speed (Fibre Channel)

Configures the transmission speed of a Fibre Channel interface.

Syntax	`speed {8 | 16 | 32 | auto}`
Parameters	Set the speed of a Fibre Channel interface to: • `8` — 8GFC • `16` — 16GFC • `32` — 32GFC • `auto` — Set the port speed to the speed of the installed media.
Defaults	Auto
Command Mode	INTERFACE
Usage Information	The `speed` command is supported only on Management and Fibre Channel interfaces. This command is not supported on Ethernet interfaces. • To configure oversubscription for bursty storage traffic on a FC interface, use the `speed` command. Oversubscription allows a port to operate faster, but may result in traffic loss. For example, QSFP28 port

groups in 4x8GFC mode support 16GFC oversubscription on member interfaces. QSFP28 breakout interfaces in 4x16GFC mode support 32GFC oversubscription.

- The `no` version of this command resets the port speed to the default value `auto`.

Example

```
OS10(conf-if-fc-1/1/2)# speed 16
```

Supported Releases 10.3.1E or later

speed (Management)

Configures the transmission speed of the Management interface.

Syntax `speed {10 | 100 | 1000 | auto}`

Parameters Set the Management port speed to:

- `10` — 10M
- `100` — 100M
- `1000` — 1000M
- `auto` — Set the port to auto-negotiate speed with a connected device.

Defaults Auto

Command Mode INTERFACE

Usage Information The `speed` command is supported only on Management and Fibre Channel interfaces. This command is not supported on Ethernet interfaces.

- When you manually configure the Management port speed, match the speed of the remote device. Dell EMC highly recommends using auto-negotiation for the Management port.
- The `no` version of this command resets the port speed to the default value `auto`.

Example

```
OS10(conf-if-ma-1/1/1)# speed auto
```

Supported Releases 10.3.0E or later

switch-port-profile

Configures a port profile on the switch. The port profile determines the available front-panel ports and breakout modes.

Syntax `switch-port-profile` *`node/unit profile`*

Parameters

- *`node/unit`* — Enter switch information. For a standalone switch, enter `1/1`.
- *`profile`* — Enter the name of a platform-specific profile.

Default `profile-1`

Command Mode CONFIGURATION

Usage Information

- S4148-ON Series port profiles:
  - `profile-1` — SFP+ 10G ports (1-24 and 31-54) and QSFP28 100G ports (25-26 and 29-30) are enabled. QSFP28 ports support 100GE and 4x10G, 4x25G, and 2x50G breakouts.

- `profile-2` — SFP+ 10G ports (1-24 and 31-50), QSFP+ 40G ports (27-28), and QSFP28 ports in 40G mode (25-26 and 29-30) are enabled.

  QSFP+ and QSFP28 ports support 40GE and 4x10G breakouts.

- `profile-3` — SFP+ 10G ports (5-24 and 31-50), QSFP+ 40G ports (27-28), and QSFP28 ports with 40G and 100G capability (25-26 and 29-30) are enabled.

  QSFP+ ports support 40GE and 4x10G breakouts. QSFP28 ports support 100GE and 4x25G breakouts with QSFP28 transceivers, and 40GE and 4x10G breakouts with QSFP+ transceivers.

- `profile-4` — SFP+ 10G ports (5-24 and 31-50), QSFP+ 40G ports (27-28), and QSFP28 ports with 40G and 100G capability (25-26 and 29-30) are enabled.

  QSFP+ ports support 40GE and 4x10G breakouts. QSFP28 ports support 100GE and 2x50G breakouts with QSFP28 transceivers, and 40GE and 4x10G breakouts with QSFP+ transceivers.

- `profile-5` — SFP+ 10G ports (1-24 and 31-54), QSFP+ 40G ports (27-28), QSFP28 ports with 40G capability (26 and 30), and QSFP28 ports with 40G and 100G capability (25 and 29) are enabled.

  QSFP+ ports support 40GE and 4x10G breakouts. QSFP28 ports 26 and 30 support 40GE and 4x10G breakouts with QSFP+ transceivers. QSFP28 ports 25 and 29 support 100GE and 4x25G breakouts with QSFP28 transceivers, and 40GE and 4x10G breakouts with QSFP+ transceivers.

- `profile-6` — SFP+ 10G ports (1-24 and 31-54), QSFP+ 40G ports (27-28), QSFP28 ports with 40G capability (26 and 30), and QSFP28 ports with 40G and 100G capability (25 and 29) are enabled.

  QSFP+ ports support 40GE and 4x10G breakouts. QSFP28 ports 26 and 30 support 40GE and 4x10G breakouts with QSFP+ transceivers. QSFP28 ports 25 and 29 support 100GE and 2x50G breakouts with QSFP28 transceivers, and 40GE and 4x10G breakouts with QSFP+ transceivers.

• S4148U-ON Port profiles:

  - `profile-1` — SFP+ unified ports (1-24), QSFP28 unified ports (25-26 and 29-30), QSFP+ Ethernet ports (27-28), and SFP+ Ethernet ports (31-54) are enabled.
    - SFP+ unified port groups operate in FC mode with 2x16GFC breakouts (ports 1 and 3) by default and support 4x8GFC. SFP+ unified ports support Ethernet 10GE mode.
    - QSFP28 unified ports 25 and 29 operate in Ethernet 100GE mode by default, and support 40GE with QSFP+ transceivers and 4x10G breakouts. QSFP28 ports 25 and 29 support 1x32GFC, 2x16GFC, and 4x8GFC in FC mode.
    - QSFP28 unified ports 26 and 30 operate in Ethernet 40GE mode by default and support 4x10G breakouts. QSFP28 ports 26 and 30 support 1x32GFC, 2x16GFC, and 4x8GFC in FC mode.
    - QSFP+ Ethernet ports operate at 40GE by default and support 4x10G breakouts.
    - SFP+ Ethernet ports operate at 10GE.
  - `profile-2` — SFP+ unified ports (1-24), QSFP28 unified ports (25-26 and 29-30), QSFP+ Ethernet ports (27-28), and SFP+ Ethernet ports (31-54) are enabled.
    - SFP+ unified ports operate in Ethernet 10GE mode by default. SFP+ unified port groups support 4x8GFC and 2x16GFC breakouts (ports 1 and 3) in FC mode.
    - QSFP28 unified ports 25 and 29 operate in Ethernet 100GE mode by default, and support 40GE with QSFP+ transceivers and 4x10G breakouts. QSFP28 ports 25 and 29 support 1x32GFC, 2x16GFC, and 4x8GFC in FC mode.
    - QSFP28 unified ports 26 and 30 operate in Ethernet 40GE mode by default and support 4x10G breakouts. QSFP28 ports 26 and 30 support 1x32GFC, 2x16GFC, and 4x8GFC in FC mode.
    - QSFP+ Ethernet ports operate at 40GE by default and support 4x10G breakouts.
    - SFP+ Ethernet ports operate at 10GE.
  - `profile-3` — SFP+ unified ports (1-24), QSFP28 unified ports (25-26 and 29-30), and SFP+ Ethernet ports (31-54) are enabled. QSFP+ Ethernet ports (27-28) are not available.
    - SFP+ unified ports operate in Ethernet 10GE mode by default. SFP+ unified port groups support 4x8GFC and 2x16GFC breakouts (ports 1 and 3) in FC mode.
    - QSFP28 unified ports operate in Ethernet 100GE mode by default and support 4x25G and 4x10G breakouts. QSFP28 ports support 2x16GFC and 4x16GFC breakouts in FC mode.
    - SFP+ Ethernet ports operate at 10GE.
  - `profile-4` — SFP+ unified ports (1-24), QSFP28 unified ports (25-26 and 29-30), and SFP+ Ethernet ports (31-54) are enabled. QSFP+ Ethernet ports (27-28) are not available.

- SFP+ unified ports operate in Ethernet 10GE mode by default. SFP+ unified ports support 4x8FC in FC mode.
- QSFP28 unified ports operate in Ethernet 100GE mode by default, and support 2x50G, 4x25G, and 4x10G breakouts. QSFP28 ports support 4x16GFC breakouts in FC mode.
- SFP+ Ethernet ports operate at 10GE.

Usage Information

- Setting a port group in 2x16GFC mode activates odd-numbered interfaces 1 and 3. A port group in 1x32GFC mode activates only interface 1.
- To display the current port profile on a switch, use the `show switch-port-profile` command.
- To change the port profile on a switch, use the `switch-port-profile` command with the desired profile, save it to the startup configuration and use the `reload` command to apply the change. The switch reboots with new port configuration. The `no` version of the command resets to the default profile. When a switch reloads with a new port profile, the startup configuration resets to system defaults, except for the switch-port profile and these configured settings:
  - Management interface 1/1/1 configuration
  - Management IPv4/IPv6 static routes
  - System hostname
  - Unified Forwarding Table (UFT) mode
  - ECMP maximum paths

  You must manually reconfigure other settings on a switch after you apply a new port profile and use the `reload` command to apply the change.

Example

```
OS10(config)# switch-port-profile 1/1 profile-1
Warning: Switch port profile will be applied only after a save and reload. All
management port
configurations will be retained but all other configurations will be wiped out
after the reload.
OS10(config)# do write memory
OS10(config)# do reload
```

Supported Releases 10.3.0E or later

switchport access vlan

Assigns access VLAN membership to a port in L2 Access or Trunk mode.

Syntax `switchport access vlan vlan-id`

Parameters `vlan vlan-id` — Enter the VLAN ID number, from 1 to 4093.

Default VLAN 1

Command Mode INTERFACE

Usage Information This command enables L2 switching for untagged traffic and assigns a port interface to default VLAN1. Use this command to change the assignment of the access VLAN that carries untagged traffic. You must create the VLAN before you can assign an access interface to it. The `no` version of this command resets access VLAN membership on a L2 access or trunk port to VLAN 1.

Example

```
OS10(conf-if-eth1/1/3)# switchport mode access
OS10(conf-if-eth1/1/3)# switchport access vlan 100
```

Supported Releases 10.2.0E or later

switchport mode

Places an interface in L2 Access or Trunk mode.

Syntax `switchport mode {access | trunk}`

Parameters

- `access` — Enables L2 switching of untagged frames on a single VLAN.
- `trunk` — Enables L2 switching of untagged frames on the access VLAN, and of tagged frames on the VLANs specified with the `switchport trunk allowed vlan` command.

Default `access`

Command Mode INTERFACE

Usage Information

- If you assign an IP address to an interface, you cannot use this command to enable L2 switching — you must first remove the IP address.
- The `access` parameter automatically adds an interface to default VLAN1 to transmit untagged traffic. Use the `switchport access vlan` command to change the access VLAN assignment.
- The `trunk` parameter configures an interface to transmit tagged VLAN traffic. You must manually configure VLAN membership for a trunk port with the `switchport trunk allowed vlan` command.
- Use the `no switchport` command to remove all L2 configurations when you configure an L3 mode interface.
- Use the `no switchport mode` command to restore a trunk port on an interface to L2 Access mode on VLAN1.

Example

```
OS10(conf-if-eth1/1/7)# switchport mode access
```

Supported Releases 10.2.0E or later

switchport trunk allowed vlan

Configures the tagged VLAN traffic that a L2 trunk interface can carry. An L2 trunk port has no tagged VLAN membership and does not transmit tagged traffic.

Syntax `switchport trunk allowed vlan` *`vlan-id-list`*

Parameters *`vlan-id-list`* — Enter the VLAN numbers of the tagged traffic that the L2 trunk port can carry. Comma-separated and hyphenated VLAN number ranges are supported.

Default None

Command Mode INTERFACE

Usage Information Use the `no` version of this command to remove the configuration.

Example

```
OS10(conf-if-eth1/1/2)# switchport trunk allowed vlan 1000
```

```
OS10(conf-if-eth1/1/2)# no switchport trunk allowed vlan 1000
```

Supported Releases 10.2.0E or later

3

Fibre Channel

OS10 switches with Fibre Channel (FC) ports operate in one of the following modes: Direct attach (F_Port), NPIV Proxy Gateway (NPG), or FIP Snooping Bridge (FSB). In the FSB mode, you cannot use the FC ports.

F_Port

Fibre Channel fabric port (F_Port) is the switch port that connects the FC fabric to a node. S4148U-ON switches support F_Port.

Enable Fibre Channel F_Port mode globally using the `feature fc domain-ID` *`domain-ID`* command in CONFIGURATION mode.

```
OS10(config)# feature fc domain-id 100
```

NPIV Proxy Gateway

A node port (N_Port) is a port on a network node that acts as a host or storage device, and is used in FC point-to-point or FC switched fabric topologies.

N_Port ID Virtualization (NPIV) allows multiple N_Port IDs to share a single physical N_Port.

The NPIV Proxy Gateway (NPG) provides Fibre Channel over Ethernet (FCoE) to Fibre Channel (FC) bridging and vice versa. Starting from OS 10.4.1, NPG supports FC to FC switching as well.

The S4148U-ON supports both, CNA and HBA, in NPG mode.

Enable NPG mode globally using the `feature fc npg` command in CONFIGURATION mode.

To change the port mode from default N_Port, use the `fc port-mode F` command on FC interfaces.

NOTE: In a switch configured in NPG or F-Port mode, OS10 does not support scale profile VLAN configuration.To use scale profile configuration in NPG or F-Port mode, enable CPU-based VLAN flooding on the vfabric VLAN using the `mode L3` command.

FIP snooping bridge

FCoE encapsulates FC frames over Ethernet networks. FCoE Initialization protocol (FIP) establishes FC connectivity with Ethernet ports. FSB implements security characteristics to admit valid FCoE traffic in the Ethernet networks. FIP and FCoE provide FC emulation-over-Ethernet links. OS10 switches with Ethernet ports operate in FSB.

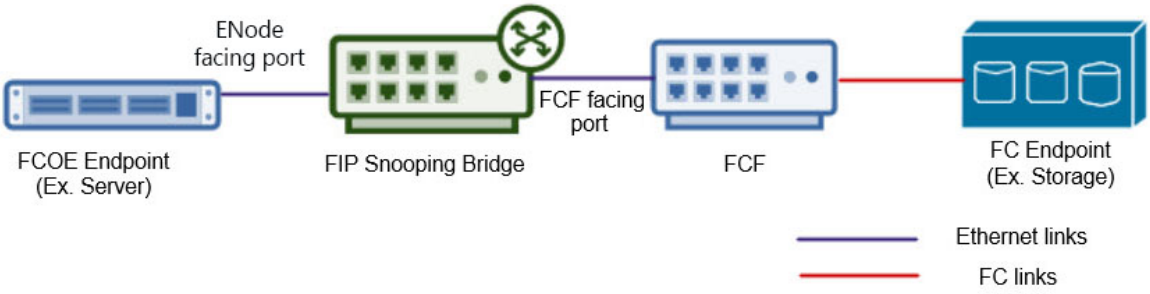
```
OS10(config)# feature fip-snooping
```

An Ethernet switch operating in FSB mode snoops FIP packets on FCoE-enabled VLANs and discovers the following information:

- End nodes (ENodes)
- Fibre Channel Forwarder (FCF)
- Connections between ENodes and FCFs
- Sessions between ENodes and FCFs

NOTE: OS10 supports multiple ENodes in F_Port mode.

Using the discovered information, the switch installs ACL entries that provide security and point-to-point link emulation.



Terminology

ENode	End Node or FCoE node
FC	Fibre Channel
FC ID	A 3-byte address used by FC to identify the end points
FC Map	A 3-byte prefix configured per VLAN, used to frame FCoE MAC address
FCF	Fibre Channel Forwarder
FCoE	Fibre Channel over Ethernet
FCoE MAC	Unique MAC address used to identify an FCoE session. This is a combination of FC ID and FC Map.
FIP	FCoE Initialization Protocol
NPG	NPIV Proxy Gateway
NPIV	N_Port ID Virtualization

Virtual fabric

Virtual fabrics (vfabric) divide a physical fabric into logical fabrics. Manage each vfabric independently. The fabric ID identifies each vfabric. You can configure only one vfabric in F_Port mode and multiple vfabrics in NPG mode. F_Port and NPG modes are mutually exclusive.

If you have already configured a vfabric in F_Port mode, while configuring vfabrics in NPG mode, disable F_Port mode. When you disable F_Port mode, the existing vfabric is removed and you must configure new vfabrics in NPG mode. If you are moving from NPG mode to F_Port mode, disable NPG mode and create the new vfabric in F_Port mode.

Zoning allows you to increase network security by partitioning the devices connected to the vfabric into subsets. Partitioning restricts unnecessary interactions between the members of vfabric. For more information, see Fibre Channel zoning.

After configuring a vfabric ID, you can create a name, associate a VLAN to carry traffic to the vfabric, configure FCoE parameters, configure the default zone, and activate the zoneset.

NOTE: Do not associate a VLAN that is already in use as a vfabric VLAN.

To configure a vfabric in F_Port mode:

1. Configure a vfabric using the `vfabric fabric-ID` command in CONFIGURATION mode. The switch enters vfabric CONFIGURATION mode
2. Associate a VLAN ID to the vfabric with the `vlan vlan-ID` command.
3. Add an FC map with the `fcoe fcmap fc-map` command.
4. Activate a zoneset using the `zoneset activate zoneset-name` command.
5. Allow access to all logged-in members in the absence of an active zoneset configuration using the `zone default-zone permit` command. The logged-in members are the FC nodes that are successfully logged into the FC fabric, identified by the vfabric.

6 (Optional) Add a name to the vfabric using the `name vfabric-name` command.

7 Apply the vfabric to FC interfaces using the `vfabric fabric-ID` command in FC INTERFACE mode.

Example configuration of vfabric in F_Port mode

```
OS10(config)# vfabric 100
OS10(conf-vfabric-100)# name 100
OS10(conf-vfabric-100)# vlan 1023
OS10(conf-vfabric-100)# fcoe fcmap 0xEFC64
OS10(conf-vfabric-100)# zoneset activate set
OS10(conf-vfabric-100)# zone default-zone permit
OS10(conf-vfabric-100)# exit
OS10(config)# interface fibrechannel 1/1/1
OS10(conf-if-fc1/1/1)# vfabric 100
```

View vfabric configuration

```
OS10(conf-vfabric-100)# show configuration
!
vfabric 100
 name 100
 vlan 1023
 fcoe fcmap 0xEFC64
 zoneset activate set
 zone default-zone permit
```

```
OS10# show vfabric
Fabric Name           100
Fabric Type           FPORT
Fabric Id             100
Vlan Id               1023
FC-MAP                0xEFC64
Config-State          ACTIVE
Oper-State            UP
=========================================
Switch Config Parameters
=========================================
Domain ID             100
=========================================
Switch Zoning Parameters
=========================================
Default Zone Mode:    Allow
Active ZoneSet:       set
=========================================
Members
fibrechannel1/1/1
fibrechannel1/1/2
fibrechannel1/1/3
fibrechannel1/1/4
fibrechannel1/1/5
fibrechannel1/1/6
fibrechannel1/1/7
fibrechannel1/1/8
fibrechannel1/1/9
fibrechannel1/1/10
fibrechannel1/1/11
fibrechannel1/1/12
fibrechannel1/1/15
fibrechannel1/1/17
fibrechannel1/1/18
fibrechannel1/1/19
fibrechannel1/1/20
fibrechannel1/1/21
fibrechannel1/1/22
fibrechannel1/1/23
fibrechannel1/1/24
fibrechannel1/1/25:1
fibrechannel1/1/29:1
fibrechannel1/1/30:1
```

```
fibrechannel1/1/30:3
==========================================
```

To configure a vfabric in NPG mode:

1 Configure a vfabric using the `vfabric fabric-ID` command in CONFIGURATION mode. The switch enters vfabric CONFIGURATION mode.

2 Associate a VLAN ID to the vfabric with the `vlan vlan-ID` command.

3 Add FCoE parameters with the `fcoe {fcmap fc-map | fcf-priority fcf-priority-value | fka-adv-period adv-period | vlan-priority vlan-priority-value | keep-alive}` command.

4 (Optional) Add a name to the vfabric using the `name vfabric-name` command.

5 Apply the vfabric to interfaces using the `vfabric fabric-ID` command in INTERFACE mode.

Configure vfabric in NPG mode

```
OS10(config)# vfabric 10
OS10(conf-vfabric-10)# name 10
OS10(conf-vfabric-10)# vlan 100
OS10(conf-vfabric-10)# fcoe fcmap 0x0efc01
OS10(conf-vfabric-10)# fcoe fcf-priority 128
OS10(conf-vfabric-10)# fcoe fka-adv-period 8
OS10(conf-vfabric-10)# fcoe vlan-priority 3
OS10(conf-vfabric-10)# exit
OS10(config)# interface ethernet 1/1/31
OS10(conf-if-eth1/1/31)# vfabric 10
```

View vfabric configuration

```
OS10(conf-vfabric-10)# show configuration
!
vfabric 10
 name 10
 vlan 100
 fcoe fcmap 0xEFC01
 fcoe fcf-priority 128
 fcoe fka-adv-period 8
 fcoe vlan-priority 3
```

```
OS10# show vfabric
Fabric Name 10
Fabric Type NPG
Fabric Id 10
Vlan Id 100
FC-MAP 0xEFC01
Vlan priority 3
FCF Priority 128
FKA-Adv-Period Enabled,8
Config-State ACTIVE
Oper-State DOWN
==========================================
Members
==========================================
```

```
OS10# show running-configuration vfabric
!
vfabric 10
 name 10
 vlan 100
 fcoe fcmap 0xEFC01
 fcoe fcf-priority 128
 fcoe fka-adv-period 8
 fcoe vlan-priority 3
```

Fibre Channel zoning

Fibre Channel (FC) zoning partitions a FC fabric into subsets to restrict unnecessary interactions, improve security, and manage the fabric more effectively. Create zones and add members to the zone. Identify a member by an FC alias, world wide name (WWN), or FC ID. A zone can have a maximum of 255 unique members. Create zonesets and add the zones to a zoneset. A switch can have multiple zonesets, but you can activate only one zoneset at a time in a fabric.

1. (Optional) Create an FC alias using the `fc alias alias-name` command in CONFIGURATION mode. The switch enters Alias CONFIGURATION mode.
2. Add members to the alias using the `member {wwn wwn-ID | fc-id fc-id}` command in Alias CONFIGURATION mode. You can add a maximum of 255 unique members.
3. Create a zone using the `fc zone zone-name` command in CONFIGURATION mode. The switch enters Zone CONFIGURATION mode.
4. Add members to the zone with the `member {alias-name alias-name | wwn wwn-ID | fc-id fc-id}` command in Zone CONFIGURATION mode.
5. Create a zoneset using the `fc zoneset zoneset-name` command in CONFIGURATION mode. The switch enters Zoneset CONFIGURATION mode.
6. Add the existing zones to the zoneset with the `member zone-name` command in Zoneset CONFIGURATION mode.
7. Activate the zoneset using the `zoneset activate zoneset-name` command in vfabric CONFIGURATION mode. The members in the zoneset become active.
8. Allow access between all the logged-in FC nodes in the absence of an active zoneset configuration using the `zone default-zone permit` command in vfabric CONFIGURATION mode. A default zone advertises a maximum of 255 members in the registered state change notification (RSCN) message.

> **NOTE: The default-zone allows or denies access to the FC nodes when an active zoneset is not available. When the default-zone action is set to `permit`, the switch allows communication between all the possible pairs of FC nodes. When you do not configure the default-zone action, the switch denies any communication between FC nodes.**

To configure the vfabric on FC interfaces, associate a VLAN ID to the vfabric and add an FC map. For more information, see Virtual fabric.

Configure FC zoning

```
OS10(config)# fc zone hba1
OS10(config-fc-zone-hba1)# member wwn 10:00:00:90:fa:b8:22:19
OS10(config-fc-zone-hba1)# member wwn 21:00:00:24:ff:7b:f5:c8
OS10(config-fc-zone-hba1)# exit

OS10(config)# fc zoneset set
OS10(conf-fc-zoneset-set)# member hba1
OS10(conf-fc-zoneset-set)# exit

OS10(config)# vfabric 100
OS10(conf-vfabric-100)# zoneset activate set
OS10(conf-vfabric-100)# zone default-zone permit
```

View FC zone configuration

```
OS10(config-fc-zone-hba1)# show configuration
!
fc zone hba1
 member wwn 21:00:00:24:ff:7b:f5:c8
 member wwn 10:00:00:90:fa:b8:22:19
```

```
OS10# show fc zone

    Zone Name                   Zone Member
================================================
hba1                        21:00:00:24:ff:7b:f5:c8
                            10:00:00:90:fa:b8:22:19

hba2                        20:01:00:0e:1e:e8:e4:99
```

```
                            50:00:d3:10:00:ec:f9:1b
                            50:00:d3:10:00:ec:f9:05
                            50:00:d3:10:00:ec:f9:1f
                            20:35:78:2b:cb:6f:65:57
```

View FC zoneset configuration

```
OS10(conf-fc-zoneset-set)# show configuration
!
fc zoneset set
 member hba1
 member hba2
```

```
OS10# show fc zoneset active

vFabric id: 100
Active Zoneset: set
ZoneName                      ZoneMember
================================================
hba2                        *20:01:00:0e:1e:e8:e4:99
                             20:35:78:2b:cb:6f:65:57
                             50:00:d3:10:00:ec:f9:05
                             50:00:d3:10:00:ec:f9:1b
                             50:00:d3:10:00:ec:f9:1f

hba1                         *10:00:00:90:fa:b8:22:19
                             *21:00:00:24:ff:7b:f5:c8
```

```
OS10# show fc zoneset set
ZoneSetName         ZoneName             ZoneMember
============================================================
set                 hba1              21:00:00:24:ff:7b:f5:c8
                                      10:00:00:90:fa:b8:22:19
                                      21:00:00:24:ff:7f:ce:ee
                                      21:00:00:24:ff:7f:ce:ef

                    hba2              20:01:00:0e:1e:e8:e4:99
                                      50:00:d3:10:00:ec:f9:1b
                                      50:00:d3:10:00:ec:f9:05
                                      50:00:d3:10:00:ec:f9:1f
                                      20:35:78:2b:cb:6f:65:57
```

F_Port on Ethernet

OS10 supports configuring F_Port mode on an Ethernet port that connects to converged network adapters (CNA). After enabling F_Port mode, configure a vfabric and apply the vfabric to Ethernet ports connected to CNA. You can configure only one vfabric in F_Port mode.

You can apply the configured vfabric to multiple Ethernet interfaces. You can also add Ethernet interfaces to a port-channel and apply the vfabric to the port-channel.

Example configuration

```
OS10(config)# feature fc domain-id 100
OS10(config)# vfabric 100
OS10(conf-vfabric-100)# name 100
OS10(conf-vfabric-100)# vlan 1023
OS10(conf-vfabric-100)# fcoe fcmap 0xEFC64
OS10(conf-vfabric-100)# zoneset activate set
OS10(conf-vfabric-100)# zone default-zone permit
OS10(conf-vfabric-100)# exit
OS10(config)# interface ethernet 1/1/30
OS10(conf-if-eth1/1/30)# vfabric 100
```

Pinning FCoE traffic to a specific port of a port-channel

You can isolate FIP and FCoE traffic by configuring a pinned port at the FCoE LAG.

FCoE LAG is the port-channel used for FIP and FCoE traffic in the intermediate switches between server and storage devices.

VLT provides Active/Active LAN connectivity on converged links by forwarding traffic in multiple paths to multiple upstream devices without STP blocking any of the uplinks. This works for Ethernet traffic, but FCoE requires dedicated links for each SAN Fabric. FCoE traffic sent on VLT breaks SAN fabric isolation.

The FC sessions form between FC nodes and FCoE sessions happen between Ethernet nodes.

To form FC or FCoE sessions, the fabric login request and reply must traverse the switch through the same port. The fabric login request initiated from the server through the switch reaches the SAN Fabric. The login accept response is hashed out to any of the ports in the port-channel. If the server receives the response on a different port than where the request was sent, the server keeps retrying the request. Because of this action, the FC or FCoE sessions learnt based on the login accept response change to the unstable state. The sessions keep flapping until the request and response converge in the same port. To avoid this, pin one of the ports in the port-channel.

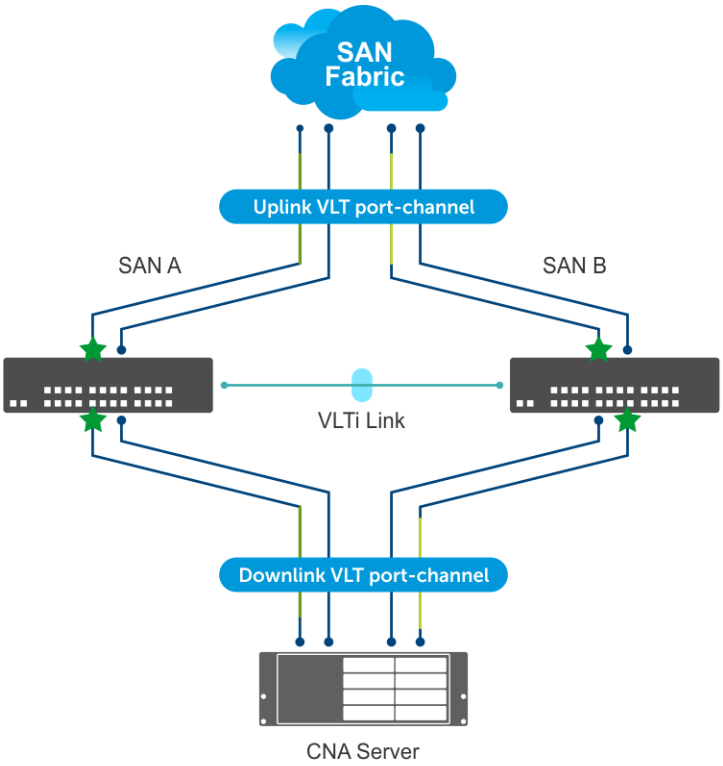
To support FCoE on multi-level VLT networks, use port pinning in FCoE LAGs. Port pinning is a static configuration that restricts the FIP and FCoE traffic to one port of the port-channel overriding hardware LAG hashing. The system classifies and redirects the packets exchanged during FCoE sessions to the port based on the ACL configuration. The remaining Ethernet traffic flows through both the pinned port and other ports in the port-channel, based on LAG hashing. Dell EMC recommends to use pinned port if there are more than one port in FCoE LAG. In a VLT network, the server has two unique FCoE sessions to SAN fabric and the traffic flows based on pinned port configuration. If there is only one port in the port-channel, there is no need for a pinned port.

NOTE: The pinned port configuration is supported on FSB, Ethernet downlink port-channel of NPG, and F_Port mode.

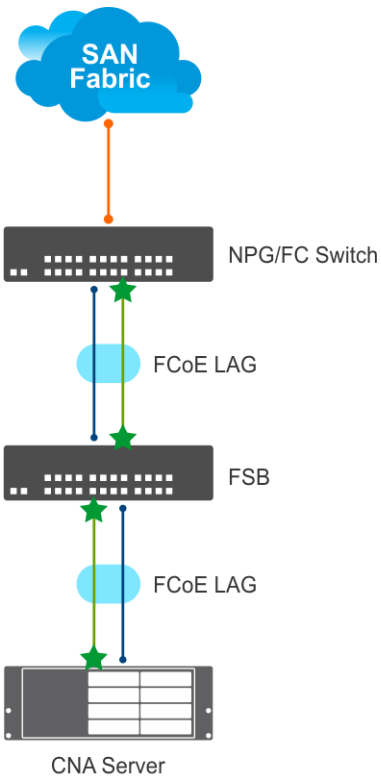
Limitations:

- The system uses an ACL table for ENode MAC with a redirect port option similar to FCF. This limits the number of FC or FCoE sessions.
- When the pinned port goes down, you must manually re-configure another active port in the port-channel as pinned port. You can perform this re-configuration only in the intermediate switches, but not in the server.
- If there is a mismatch in the configuration or if the pinned port goes down, the system does not use other ports in port-channel even if there is a valid path to server and storage device.
- When you add or remove a pinned port when FCoE sessions are active, the system clears and re-initiates the FCoE sessions based on the configuration. The system displays warning messages during the configuration.

The following illustrations show VLT and non-VLT networks with FCoE traffic flowing through pinned port.


SAN Fabric
Uplink VLT port-channel
SAN A
SAN B
VLTi Link
Downlink VLT port-channel
CNA Server


Ethernet
Converged
FCoE Session - SAN A
FCoE Session - SAN B
Pinned Port


SAN Fabric
NPG/FC Switch
FCoE LAG
FSB
FCoE LAG
CNA Server


Fiber Channel
Converged
FCoE Session
Pinned Port

Sample FSB configuration on VLT network

1 Enable the FIP snooping feature globally.

```
OS10(config)# feature fip-snooping
```

2 Create the FCoE VLAN.

```
OS10(config)#interface vlan 1001
OS10(conf-if-vl-1001)# fip-snooping enable
```

3 Configure the VLTi interface.

```
OS10(config)# interface ethernet 1/1/27
OS10(conf-if-eth1/1/27)# no shutdown
OS10(conf-if-eth1/1/27)# no switchport
```

4 Configure the VLT.

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination 10.16.151.110  <<Enter the management IP address of the
VLT peer>>
OS10(conf-vlt-1)# discovery-interface ethernet1/1/27
```

5 Enable DCBX.

```
OS10(config)# dcbx enable
```

6 Enable the PFC parameters on the interfaces.

```
OS10(config)# class-map type network-qos fcoematch
OS10(config-cmap-nqos)# match qos-group 3
OS10(config-cmap-nqos)# exit
OS10(config)# policy-map type network-qos PFC
OS10(config-pmap-network-qos)# class fcoematch
OS10(config-pmap-c-nqos)# pause
OS10(config-pmap-c-nqos)# pfc-cos 3
```

7 Create uplink and downlink port-channels, and configure the FCF facing port.

```
OS10(config)# interface port-channel 10
OS10(conf-if-po-10)# description uplink_VLT_LAG
OS10(conf-if-po-10)# no shutdown
OS10(conf-if-po-10)# switchport mode trunk
OS10(conf-if-po-10)# switchport access vlan 1
OS10(conf-if-po-10)# switchport trunk allowed vlan 1001,10
OS10(conf-if-po-10)# vlt-port-channel 1
OS10(conf-if-po-10)# fip-snooping port-mode fcf
```

```
OS10(config)# interface port-channel 20
OS10(conf-if-po-20)# description downlink_VLT_LAG
OS10(conf-if-po-20)# no shutdown
OS10(conf-if-po-20)# switchport mode trunk
OS10(conf-if-po-20)# switchport access vlan 1
OS10(conf-if-po-20)# switchport trunk allowed vlan 1001,10
OS10(conf-if-po-20)# vlt-port-channel 2
```

8 Apply the PFC configuration on downlink and uplink interfaces. In addition, include the interfaces to the port-channel and configure one of the interfaces as pinned-port.

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# description uplink_port_channel_member1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# channel-group 10 mode active
OS10(conf-if-eth1/1/1)# fcoe-pinned-port
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# service-policy input type network-qos PFC
OS10(conf-if-eth1/1/1)# priority-flow-control mode on
```

```
OS10(config)# interface ethernet 1/1/2
OS10(conf-if-eth1/1/2)# description uplink_port_channel_member2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# channel-group 10 mode active
OS10(conf-if-eth1/1/2)# no switchport
```

```
OS10(conf-if-eth1/1/2)# service-policy input type network-qos PFC
OS10(conf-if-eth1/1/2)# priority-flow-control mode on
```

```
OS10(config)# interface ethernet 1/1/3
OS10(conf-if-eth1/1/3)# description downlink_port_channel_member1
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# channel-group 20 mode active
OS10(conf-if-eth1/1/3)# fcoe-pinned-port
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# service-policy input type network-qos PFC
OS10(conf-if-eth1/1/3)# priority-flow-control mode on
```

```
OS10(config)# interface ethernet 1/1/4
OS10(conf-if-eth1/1/4)# description downlink_port_channel_member2
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# channel-group 20 mode active
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# service-policy input type network-qos PFC
OS10(conf-if-eth1/1/4)# priority-flow-control mode on
```

View the configuration

VLT details:

```
OS10# show vlt 1
Domain ID                  : 1
Unit ID                    : 2
Role                       : primary
Version                    : 2.0
Local System MAC address   : 50:9a:4c:d3:cf:70
Primary priority           : 32768
VLT MAC address            : 50:9a:4c:d3:cf:70
IP address                 : fda5:74c8:b79e:1::2
Delay-Restore timer        : 90 seconds
Peer-Routing               : Disabled
Peer-Routing-Timeout timer : 0 seconds
VLTi Link Status
   port-channel1000       : up

VLT Peer Unit ID    System MAC Address    Status    IP Address              Version
--------------------------------------------------------------------------------
  1                 50:9a:4c:d3:e2:f0      up       fda5:74c8:b79e:1::1     2.0
```

```
OS10# show vlt 1 vlt-port-detail
vlt-port-channel ID : 1
VLT Unit ID    Port-Channel      Status    Configured ports    Active ports
--------------------------------------------------------------------------
  1            port-channel10      up        2                   2
* 2            port-channel10      up        2                   2
vlt-port-channel ID : 2
VLT Unit ID    Port-Channel      Status    Configured ports    Active ports
--------------------------------------------------------------------------
  1            port-channel20      up        2                   2
* 2            port-channel20      up        2                   2
```

Discovered ENodes:

```
OS10# show fcoe enode
Enode MAC           Enode Interface    VLAN   FCFs    Sessions
------------------------------------------------------------------
f4:e9:d4:a4:7d:c3   Po 20(Eth 1/1/3)   1001      1         1
```

Discovered FCFs:

```
OS10# show fcoe fcf
FCF MAC             FCF Interface    VLAN     FC-MAP    FKA_ADV_PERIOD    No. of Enodes
-------------------------------------------------------------------------------------
14:18:77:20:78:e0    Po 10(Eth 1/1/1) 1001   0e:fc:00       8000               1
```

FCoE sessions:

```
Enode MAC            Enode Interface   FCF MAC            FCF interface  VLAN   FCoE
MAC           FC-ID   PORT WWPN                 PORT WWNN
--------------------------------------------------------------------------------------------
------------------------------------------------------------
f4:e9:d4:a4:7d:c3   Po20(Eth 1/1/3)         14:18:77:20:78:e0  Po 10(Eth 1/1/1)
1001    0e:fc:00:01:00:00  01:34:02 20:01:f4:e9:d4:a4:7d:c3  20:00:f4:e9:d4:a4:7d:c3
```

Pinned port status:

```
OS10# show fcoe pinned-port
Interface           pinned-port        FCoE Status
----------------- ---------------- -----------------
Po 10               Eth 1/1/1          Up
Po 20               Eth 1/1/3          Up
```

Sample FC Switch configuration on VLT network

1 Enable the F_PORT mode.

```
OS10(config)# feature fc domain-id 1
```

2 Create the FC zones.

```
OS10(config)# fc zone zoneA
OS10(config-fc-zone-zoneA)# member wwn 10:00:00:90:fa:b8:22:19 <<Enter the WWN of Initiator
CNA>>
OS10(config-fc-zone-zoneA)# member wwn 21:00:00:24:ff:7b:f5:c8 <<Enter the WWN of Target>>
```

3 Create the FC zoneset.

```
OS10(config)# fc zoneset zonesetA
OS10(conf-fc-zoneset-zonesetA)# member zoneA
```

4 Create the vfabric VLAN.

```
OS10(config)# interface vlan 1001
```

5 Create vfabric and activate the FC zoneset.

```
OS10(config)# vfabric 1
OS10(conf-vfabric-1)# vlan 1001
OS10(conf-vfabric-1)# fcoe fcmap 0xEFC00
OS10(conf-vfabric-1)# zoneset activate zonesetA
```

6 Configure the VLTi interface.

```
OS10(config)# interface ethernet 1/1/27
OS10(conf-if-eth1/1/27)# no shutdown
OS10(conf-if-eth1/1/27)# no switchport
```

7 Configure the VLT.

```
OS10(config)# vlt-domain 10
OS10(conf-vlt-10)# backup destination 10.16.151.110
OS10(conf-vlt-10)# discovery-interface ethernet1/1/27
```

8 Enable DCBX.

```
OS10(config)# dcbx enable
```

9 Apply the vfabric on the interfaces.

```
OS10(config)# interface port-channel 10
OS10(conf-if-po-10)# description downlink_VLT_LAG_to FSB
OS10(conf-if-po-10)# no shutdown
OS10(conf-if-po-10)# switchport mode trunk
OS10(conf-if-po-10)# switchport access vlan 1
OS10(conf-if-po-10)# switchport trunk allowed vlan 10
OS10(conf-if-po-10)# vlt-port-channel 1
OS10(conf-if-po-10)# vfabric 1
```

```
OS10(config)# interface fibrechannel 1/1/26
OS10(conf-if-fc1/1/26)# description target_connected_port
OS10(conf-if-fc1/1/26)# no shutdown
OS10(conf-if-fc1/1/26)# vfabric 1
```

10 Apply the PFC configuration on the downlink interfaces. Include the interfaces to the port-channel and configure one of the interfaces as pinned-port.

```
OS10(config)# interface ethernet 1/1/9
OS10(conf-if-eth1/1/9)# description downlink_port_channel_member1
OS10(conf-if-eth1/1/9)# no shutdown
OS10(conf-if-eth1/1/9)# channel-group 10 mode active
OS10(conf-if-eth1/1/9)# fcoe-pinned-por
OS10(conf-if-eth1/1/9)# no switchport
OS10(conf-if-eth1/1/9)# service-policy input type network-qos PFC
OS10(conf-if-eth1/1/9)# priority-flow-control mode on
```

```
OS10(config)# interface ethernet 1/1/10
OS10(conf-if-eth1/1/10)# description downlink_port_channel_member2
OS10(conf-if-eth1/1/10)# no shutdown
OS10(conf-if-eth1/1/10)# channel-group 10 mode active
OS10(conf-if-eth1/1/10)# no switchport
OS10(conf-if-eth1/1/10)# service-policy input type network-qos PFC
OS10(conf-if-eth1/1/10)# priority-flow-control mode on
```

View configuration

Name server entries:

```
OS10# show fc ns switch brief
Total number of devices = 2
Intf#                    Domain    FC-ID            Enode-WWPN            Enode-WWNN
port-channel10(Eth 1/1/9) 1      01:00:00        20:01:f4:e9:d4:a4:7d:c3
20:00:f4:e9:d4:a4:7d:c3
fibrechannel1/1/26        1      01:68:00        21:00:00:24:ff:7c:ae:0e  21:00:00:24:ff:7c:ae:
0e
```

Zoneset details:

```
vFabric id: 1
Active Zoneset: zonesetA
ZoneName                         ZoneMember
============================================================
zoneA                            *20:01:f4:e9:d4:a4:7d:c3
                                 *21:00:00:24:ff:7c:ae:0e
```

Pinned port status:

```
OS10# show fcoe pinned-port
Interface             pinned-port         FCoE Status
----------------- ---------------- -----------------
Po 10                  Eth 1/1/9                Up
```

Sample FSB configuration on non-VLT network

The following examples illustrate configurations in intermediate switches in non-vlt network, to communicate with server.

1 Enable the FIP snooping feature globally.

```
OS10(config)# feature fip-snooping
```

2 Create the FCoE VLAN.

```
OS10(config)#interface vlan 1001
OS10(conf-if-vl-1001)# fip-snooping enable
```

3 Enable DCBX.

```
OS10(config)# dcbx enable
```

4 Enable the PFC parameters on the interfaces.

```
OS10(config)# class-map type network-qos fcoematch
OS10(config-cmap-nqos)# match qos-group 3
OS10(config-cmap-nqos)# exit
OS10(config)# policy-map type network-qos PFC
OS10(config-pmap-network-qos)# class fcoematch
```

```
OS10(config-pmap-c-nqos)# pause
OS10(config-pmap-c-nqos)# pfc-cos 3
```

5 Create uplink and downlink port-channels, and configure the FCF facing port.

```
OS10(config)# interface port-channel 10
OS10(conf-if-po-10)# no shutdown
OS10(conf-if-po-10)# switchport mode trunk
OS10(conf-if-po-10)# switchport access vlan 1
OS10(conf-if-po-10)# switchport trunk allowed vlan 1001,10
OS10(conf-if-po-10)# fip-snooping port-mode fcf
```

```
OS10(config)# interface port-channel 20
OS10(conf-if-po-20)# no shutdown
OS10(conf-if-po-20)# switchport mode trunk
OS10(conf-if-po-20)# switchport access vlan 1
OS10(conf-if-po-20)# switchport trunk allowed vlan 1001,10
```

6 Apply the PFC configuration on downlink and uplink interfaces. In addition, include the interfaces to the port-channel and configure one of the interfaces as pinned-port.

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# channel-group 10 mode active
OS10(conf-if-eth1/1/1)# fcoe-pinned-port
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# service-policy input type network-qos PFC
OS10(conf-if-eth1/1/1)# priority-flow-control mode on
```

```
OS10(config)# interface ethernet 1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# channel-group 10 mode active
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/2)# service-policy input type network-qos PFC
OS10(conf-if-eth1/1/2)# priority-flow-control mode on
```

```
OS10(config)# interface ethernet 1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# channel-group 20 mode active
OS10(conf-if-eth1/1/3)# fcoe-pinned-port
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# service-policy input type network-qos PFC
OS10(conf-if-eth1/1/3)# priority-flow-control mode on
```

```
OS10(config)# interface ethernet 1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# channel-group 20 mode active
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# service-policy input type network-qos PFC
OS10(conf-if-eth1/1/4)# priority-flow-control mode on
```

View the configuration

Discovered ENodes:

```
OS10# show fcoe enode
Enode MAC           Enode Interface    VLAN   FCFs    Sessions
--------------------------------------------------------------
f4:e9:d4:a4:7d:c3   Po 20(Eth 1/1/3)   1001      1         1
```

Discovered FCFs:

```
OS10# show fcoe fcf
FCF MAC             FCF Interface    VLAN     FC-MAP    FKA_ADV_PERIOD    No. of Enodes
-------------------------------------------------------------------------------------
14:18:77:20:78:e0    Po 10(Eth 1/1/1) 1001   0e:fc:00       8000                1
```

FCoE sessions:

```
Enode MAC           Enode Interface   FCF MAC            FCF interface  VLAN   FCoE
MAC           FC-ID   PORT WWPN                PORT WWNN
-----------------------------------------------------------------------------------------------
```

```
-------------------------------------------------------------
f4:e9:d4:a4:7d:c3   Po20(Eth 1/1/3)         14:18:77:20:78:e0  Po 10(Eth 1/1/1)
1001    0e:fc:00:01:00:00  01:34:02 20:01:f4:e9:d4:a4:7d:c3  20:00:f4:e9:d4:a4:7d:c3
```

Pinned port status:

```
OS10# show fcoe pinned-port
Interface           pinned-port        FCoE Status
----------------- ---------------- -----------------
Po 10               Eth 1/1/1          Up
Po 20               Eth 1/1/3          Up
```

Sample FC Switch configuration on non-VLT network

1 Enable the F_PORT mode.

```
OS10(config)# feature fc domain-id 1
```

2 Create the FC zones.

```
OS10(config)# fc zone zoneA
OS10(config-fc-zone-zoneA)# member wwn 10:00:00:90:fa:b8:22:19 <<Enter the WWN of Initiator
CNA>>
OS10(config-fc-zone-zoneA)# member wwn 21:00:00:24:ff:7b:f5:c8 <<Enter the WWN of Target>>
```

3 Create the FC zoneset.

```
OS10(config)# fc zoneset zonesetA
OS10(conf-fc-zoneset-zonesetA)# member zoneA
```

4 Create the vfabric VLAN.

```
OS10(config)# interface vlan 1001
```

5 Create vfabric and activate the FC zoneset.

```
OS10(config)# vfabric 1
OS10(conf-vfabric-1)# vlan 1001
OS10(conf-vfabric-1)# fcoe fcmap 0xEFC00
OS10(conf-vfabric-1)# zoneset activate zonesetA
```

6 Enable DCBX.

```
OS10(config)# dcbx enable
```

7 Apply the vfabric on the interfaces.

```
OS10(config)# interface port-channel 10
OS10(conf-if-po-10)# no shutdown
OS10(conf-if-po-10)# switchport mode trunk
OS10(conf-if-po-10)# switchport access vlan 1
OS10(conf-if-po-10)# switchport trunk allowed vlan 10
OS10(conf-if-po-10)# vfabric 1
```

```
OS10(config)# interface fibrechannel 1/1/26
OS10(conf-if-fc1/1/26)# description target_connected_port
OS10(conf-if-fc1/1/26)# no shutdown
OS10(conf-if-fc1/1/26)# vfabric 1
```

8 Apply the PFC configuration on the downlink interfaces. Include the interfaces to the port-channel and configure one of the interfaces as pinned-port.

```
OS10(config)# interface ethernet 1/1/9
OS10(conf-if-eth1/1/9)# no shutdown
OS10(conf-if-eth1/1/9)# channel-group 10 mode active
OS10(conf-if-eth1/1/9)# fcoe-pinned-por
OS10(conf-if-eth1/1/9)# no switchport
OS10(conf-if-eth1/1/9)# service-policy input type network-qos PFC
OS10(conf-if-eth1/1/9)# priority-flow-control mode on
```

```
OS10(config)# interface ethernet 1/1/10
OS10(conf-if-eth1/1/10)# no shutdown
OS10(conf-if-eth1/1/10)# channel-group 10 mode active
OS10(conf-if-eth1/1/10)# no switchport
OS10(conf-if-eth1/1/10)# service-policy input type network-qos PFC
OS10(conf-if-eth1/1/10)# priority-flow-control mode on
```

View configuration

Name server entries:

```
OS10# show fc ns switch brief
Total number of devices = 2
Intf#                    Domain    FC-ID            Enode-WWPN             Enode-WWNN
port-channel10(Eth 1/1/9) 1      01:00:00        20:01:f4:e9:d4:a4:7d:c3
20:00:f4:e9:d4:a4:7d:c3
fibrechannel1/1/26        1      01:68:00        21:00:00:24:ff:7c:ae:0e  21:00:00:24:ff:7c:ae:
0e
```

Zoneset details:

```
vFabric id: 1
Active Zoneset: zonesetA
ZoneName                          ZoneMember
============================================================
zoneA                             *20:01:f4:e9:d4:a4:7d:c3
                                  *21:00:00:24:ff:7c:ae:0e
```

Pinned port status:

```
OS10# show fcoe pinned-port
Interface             pinned-port         FCoE Status
----------------- ---------------- -----------------
Po 10                  Eth 1/1/9                Up
```

Multi-hop FIP-snooping bridge

In typical deployments, ENode-connected switches are not directly connected to the core FC switch. Multiple intermediate switches are connected in between the switches. To establish a point-to-point connection and for secure transmission between the ENode and the FCF, all intermediate switches must support FSB to pass the FIP and FCoE traffic.

OS10 switches support the multi-hop FIP-snooping bridge. You can interconnect multiple FSBs to communicate with an upstream FC switch.

- Access FSB— This is the node that is directly connected to ENode. In the following example, FSB1 is the access FSB.
- Core FSB— This is the node that is directly connected to the FCF. In the following example, FSB2 is the core FSB.

The default port mode is the ENode. You must explicitly configure the other modes using the `fip-snooping port-mode` command. The following port modes are supported:

- ENode—Only one ENode MAC address per interface can be learnt. Configure this mode on the port connected to the ENode.
- FCF—If you configure the FSB with FCF port mode, all the FIP packets sent between the ENode and the FCF are snooped and the sessions and ENodes are learnt. Configure the FCF mode on the access FSB ports connected to the FCF-facing side.
- ENode-transit—This mode is configured on the intermediate FSBs or Layer 2 (L2) DCBX switches to which ENodes are connected.
- FCF-transit—Only the FCF advertisement and VLAN responses are snooped to learn the FCF. The FCF-transit does not learn the ENodes and session information. Configure the FCF-transit mode on the FCF-facing side of the core FSB switch.

The FCF can be in NPG or F-Port mode. The access FSB switches validate the frames and installs ACLs per the FCF to allow only FCoE and FIP traffic across the FCF.

NOTE: Port-pinning is not supported on ENodes connected to an FSB switch that is in FCF-transit mode. You cannot view the ENodes or session information using the `show` commands.

Clear virtual link frames

When an FSB clears an FCoE session for some reason, the other devices in the network, such as the ENode, FCF, and transit switches, are not informed and considers the session to be intact. FSB drops the FCoE data corresponding to the cleared session. The ENode takes a

long time to identify the issue and to recover from it. At times, interface flapping occurs and might require manual intervention to recover. To recover automatically, FSB sends a Clear Virtual Link (CVL) frame from the FCF to the ENode.

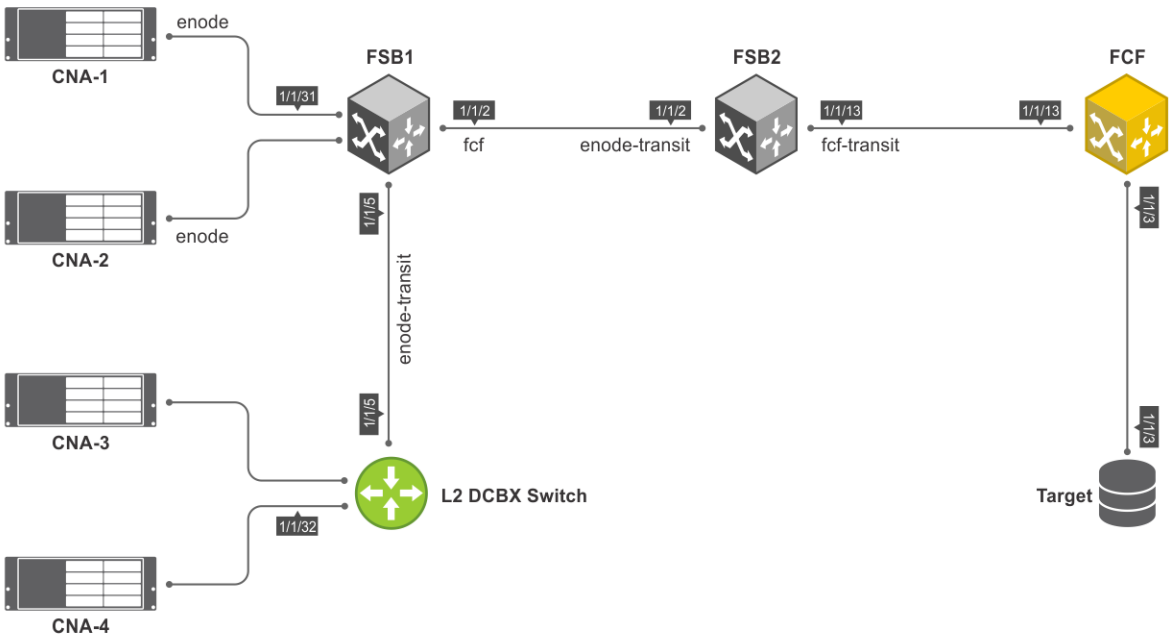
Configuration notes

- If you configure FSB with port pinning on the uplink or downlink side, you must configure the FCF-facing interface as FCF port mode.
- OS10 currently does not support a topology where a single port must be configured as both an ENode and an FCF-related port mode.
- If you configure FCF-transit port mode on an FSB, Dell EMC recommends not directly connecting ENodes on it.

Configure multi-hop FSB

The following example shows a simple multi-hop FSB setup. CNA-2 and CNA-3 shown in this topology are for illustrative purposes only. The following example does not include CNA-2 and CNA-3 configurations.

Ensure that the access and core FSB switches are running in FSB mode.



To configure multi-hop FSB:

1 Configure the L2 switch.

a Disable flow control on the interfaces connected to CNA-4 and FSB1.

```
L2switch(config)# interface ethernet 1/1/32
L2switch(conf-if-eth1/1/32)# no flowcontrol receive
L2switch(conf-if-eth1/1/32)# no flowcontrol transmit

L2switch(config)# interface ethernet 1/1/5
L2switch(conf-if-eth1/1/5)# no flowcontrol receive
L2switch(conf-if-eth1/1/5)# no flowcontrol transmit
```

b Enable DCBX.

```
L2switch(config)# dcbx enable
```

c Create a VLAN for FCoE traffic to pass through.

```
L2switch(config)# interface vlan 777
```

d Create a class-map.

```
L2switch(config)# class-map type network-qos c3
L2switch(config-cmap-nqos)# match qos-group 3
```

e Create a policy-map.

```
L2switch# configure terminal
L2switch(config)# policy-map type network-qos nqpolicy
L2switch(config-pmap-network-qos)# class c3
L2switch(config-pmap-c-nqos)# pause
L2switch(config-pmap-c-nqos)# pfc-cos 3
```

f Apply the policy-map on CNA-4 and FSB1 connected interfaces.

```
L2switch(config)# interface ethernet 1/1/32
L2switch(conf-if-eth1/1/32)# priority-flow-control mode on
L2switch(conf-if-eth1/1/32)# service-policy input type network-qos nqpolicy

L2switch(config)# interface ethernet 1/1/5
L2switch(conf-if-eth1/1/5)# priority-flow-control mode on
L2switch(conf-if-eth1/1/5)# service-policy input type network-qos nqpolicy
```

g Configure VLAN on CNA-4 and FSB1 connected interfaces.

```
L2switch(config)# interface ethernet 1/1/32
L2switch(conf-if-eth1/1/32)# switchport mode trunk
L2switch(conf-if-eth1/1/32)# switchport trunk allowed vlan 777

L2switch(config)# interface ethernet 1/1/5
L2switch(conf-if-eth1/1/5)# switchport mode trunk
L2switch(conf-if-eth1/1/5)# switchport trunk allowed vlan 777
```

2 Configure the access FSB, FSB1. This example describes

a Disable flow control on the interfaces connected to CNA1, L2 switch, and FSB2.

```
FSB1(config)# interface ethernet 1/1/31
FSB1(conf-if-eth1/1/31)# no flowcontrol receive
FSB1(conf-if-eth1/1/31)# no flowcontrol transmit

FSB1(config)# interface ethernet 1/1/5
FSB1(conf-if-eth1/1/5)# no flowcontrol receive
FSB1(conf-if-eth1/1/5)# no flowcontrol transmit

FSB1(config)# interface ethernet 1/1/2
FSB1(conf-if-eth1/1/2)# no flowcontrol receive
FSB1(conf-if-eth1/1/2)# no flowcontrol transmit
```

b Enable FIP snooping with cvl option.

```
FSB1(config)# feature fip-snooping with-cvl
```

c Enable DCBX.

```
FSB1(config)# dcbx enable
```

d Create an FCoE VLAN and configure FIP snooping on the FCoE VLAN.

```
FSB1(config)# interface vlan 777
FSB1(conf-if-vl-777)# fip-snooping enable
```

e Create a class-map.

```
FSB1(config)# class-map type network-qos c3
FSB1(config-cmap-nqos)# match qos-group 3
```

f Create a policy-map.

```
FSB1(config)# policy-map type network-qos nqpolicy
FSB1(config-pmap-network-qos)# class c3
FSB1(config-pmap-c-nqos)# pause
FSB1(config-pmap-c-nqos)# pfc-cos 3
```

g Apply the policy-map on CNA1, L2 switch, and FSB2 connected interfaces.

```
FSB1(config)# interface ethernet 1/1/31
FSB1(conf-if-eth1/1/31)# priority-flow-control mode on
FSB1(conf-if-eth1/1/31)# service-policy input type network-qos nqpolicy

FSB1(config)# interface ethernet 1/1/5
FSB1(conf-if-eth1/1/5)# priority-flow-control mode on
```

```
FSB1(conf-if-eth1/1/5)# service-policy input type network-qos nqpolicy

FSB1(config)# interface ethernet 1/1/2
FSB1(conf-if-eth1/1/2)# priority-flow-control mode on
FSB1(conf-if-eth1/1/2)# service-policy input type network-qos nqpolicy
```

h Configure VLAN on CNA1, L2 switch, and FSB2 connected interfaces.

```
FSB1(config)# interface ethernet 1/1/31
FSB1(conf-if-eth1/1/31)# switchport mode trunk
FSB1(conf-if-eth1/1/31)# switchport trunk allowed vlan 777

FSB1(config)# interface ethernet 1/1/5
FSB1(conf-if-eth1/1/5)# switchport mode trunk
FSB1(conf-if-eth1/1/5)# switchport trunk allowed vlan 777

FSB1(config)# interface ethernet 1/1/2
FSB1(conf-if-eth1/1/2)# switchport mode trunk
FSB1(conf-if-eth1/1/2)# switchport trunk allowed vlan 777
```

i Configure FIP snooping port mode on the L2 DCBX switch connected interface and FSB2 connected interface. The default port mode is ENode. Hence, CNA1-connected interface does not require additional configuration.

On the L2 DCBX switch-connected interface:

```
FSB1(config)# interface ethernet 1/1/5
FSB1(conf-if-eth1/1/5)# fip-snooping port-mode enode-transit
```

On the FSB-connected interfaces:

```
FSB1(config)# interface ethernet 1/1/2
FSB1(conf-if-eth1/1/2)# fip-snooping port-mode fcf
```

3 Configure the core FSB, FSB2.

a Disable flow control on the interfaces connected to FSB1 and FCF.

```
FSB2(config)# interface ethernet 1/1/2
FSB2(conf-if-eth1/1/2)# no flowcontrol receive
FSB2(conf-if-eth1/1/2)# no flowcontrol transmit

FSB2(config)# interface ethernet 1/1/13
FSB2(conf-if-eth1/1/13)# no flowcontrol receive
FSB2(conf-if-eth1/1/13)# no flowcontrol transmit
```

b Enable FIP snooping with cvl option.

```
FSB2(config)# feature fip-snooping with-cvl
```

c Enable DCBX.

```
FSB2(config)# dcbx enable
```

d Create an FCoE VLAN and configure FIP snooping on the FCoE VLAN.

```
FSB2(config)# interface vlan 777
FSB2(conf-if-vl-777)# fip-snooping enable
```

e Create a class-map.

```
FSB2(config)# class-map type network-qos c3
FSB2(config-cmap-nqos)# match qos-group 3
```

f Create a policy-map.

```
FSB2(config)# policy-map type network-qos nqpolicy
FSB2(config-pmap-network-qos)# class c3
FSB2(config-pmap-c-nqos)# pause
FSB2(config-pmap-c-nqos)# pfc-cos 3
```

g Apply the policy-map on FSB1 and FCF connected interfaces.

```
FSB2(config)# interface ethernet 1/1/2
FSB2(conf-if-eth1/1/2)# priority-flow-control mode on
FSB2(conf-if-eth1/1/2)# service-policy input type network-qos nqpolicy

FSB2(config)# interface ethernet 1/1/13
FSB2(conf-if-eth1/1/13)# priority-flow-control mode on
FSB2(conf-if-eth1/1/13)# service-policy input type network-qos nqpolicy
```

h Configure VLAN on FSB1 and FCF connected interfaces.

```
FSB2(config)# interface ethernet 1/1/2
FSB2(conf-if-eth1/1/2)# switchport mode trunk
FSB2(conf-if-eth1/1/2)# switchport trunk allowed vlan 777

FSB2(config)# interface ethernet 1/1/13
FSB2(conf-if-eth1/1/13)# switchport mode trunk
FSB2(conf-if-eth1/1/13)# switchport trunk allowed vlan 777
```

i Configure FIP snooping port mode on FSB1 and FCF connected interfaces.

On the FSB1-connected interface:

```
FSB2(config)# interface ethernet 1/1/2
FSB2(conf-if-eth1/1/2)# fip-snooping port-mode enode-transit
```

On the FCF-connected interface:

```
FSB2(config)# interface ethernet 1/1/13
FSB2(conf-if-eth1/1/13)# fip-snooping port-mode fcf-transit
```

4 Configure the FCF. The following configuration assumes that the FCF is in F-Port mode.

a Disable flow control on the interface connected to FSB2.

```
FCF(config)# interface ethernet 1/1/13
FCF(conf-if-eth1/1/13)# no flowcontrol receive
FCF(conf-if-eth1/1/13)# no flowcontrol transmit
```

b Enable Fiber Channel F-Port mode globally.

```
FCF(config)# feature fc domain-id 2
```

c Create zones.

```
FCF(config)# fc zone zoneA
FCF(config-fc-zone-zoneA)# member wwn 20:01:f4:e9:d4:a4:7d:c3
FCF(config-fc-zone-zoneA)# member wwn 21:00:00:24:ff:7c:ae:0e
```

d Create zoneset.

```
FCF(config)# fc zoneset zonesetA
FCF(conf-fc-zoneset-set)# member zoneA
```

e Create a vfabric VLAN.

```
FCF(config)# interface vlan 777
```

f Create vfabric and activate the zoneset.

```
FCF(config)# vfabric 2
FCF(conf-vfabric-2)# vlan 777
FCF(conf-vfabric-2)# fcoe fcmap 0xEFC00
FCF(conf-vfabric-2)# zoneset activate zonesetA
```

g Enable DCBX.

```
FCF(config)# dcbx enable
```

h Enable PFC parameters.

```
FCF(config)# class-map type network-qos c3
FCF(config-cmap-nqos)# match qos-group 3
FCF(config-cmap-nqos)# exit
FCF(config)# policy-map type network-qos nqpolicy
FCF(config-pmap-network-qos)# class c3
FCF(config-pmap-c-nqos)# pause
FCF(config-pmap-c-nqos)# pfc-cos 3
```

i Apply vfabric on FSB2 and target connected interfaces.

```
FCF(config)# interface ethernet 1/1/13
FCF(conf-if-eth1/1/13)# no shutdown
FCF(conf-if-eth1/1/13)# switchport access vlan 1
FCF(conf-if-eth1/1/13)# vfabric 2

FCF(config)# interface fibrechannel 1/1/3
FCF(conf-if-fc1/1/3)# description target_connected_port
FCF(conf-if-fc1/1/3)# no shutdown
FCF(conf-if-fc1/1/3)# vfabric 2
```

j Apply PFC on the interface connected to FSB2.

```
FCF(config)# interface ethernet 1/1/13
FCF(conf-if-eth1/1/13)# no shutdown
FCF(conf-if-eth1/1/13)# priority-flow-control mode on
FCF(conf-if-eth1/1/13)# service-policy input type network-qos nqpolicy
FCF(conf-if-eth1/1/13)# vfabric 2
```

Verify multi-hop FSB configuration

Verify the configuration using the following `show` commands:

- To verify FSB mode and the CVL status, use the `show fcoe system` command.

```
FSB1# show fcoe system
Mode                          : FSB
CVL Status                    : Enabled
FCOE VLAN List (Operational) : 777
FCFs                          : 1
Enodes                        : 1
Sessions                      : 1
```

- To verify the discovered ENodes, use the `show fcoe enode` command.

```
FSB1# show fcoe enode
Enode MAC           Enode Interface     VLAN     FCFs    Sessions
--------------------------------------------------------------
f4:e9:d4:f9:fc:40   Eth 1/1/5            777      1        1
```

- To verify the discovered FCFs, use the `show fcoe fcf` command.

```
FSB1# show fcoe fcf
FCF MAC             FCF Interface       VLAN     FC-MAP         FKA_ADV_PERIOD    No. of
Enodes      FCF Mode
-------------------------------------------------------------------------------------------
----------------
14:18:77:20:86:ce   Eth 1/1/2            777      0e:fc:00       8000
1                   F
```

```
FSB2# show fcoe fcf
FCF MAC            FCF Interface       VLAN     FC-MAP         FKA_ADV_PERIOD     No. of
Enodes      FCF Mode
-------------------------------------------------------------------------------------------
----------------
14:18:77:20:86:ce   Eth 1/1/13          777      0e:fc:00       8000
0                   FT
```

- To verify the list of FCoE sessions, use the `show fcoe sessions` command.

```
FSB1# show fcoe sessions
Enode MAC          Enode Interface  FCF MAC           FCF interface  VLAN  FCoE MAC
FC-ID    PORT WWPN
PORT WWNN
-------------------------------------------------------------------------------------------
-----------------
f4:e9:d4:f9:fc:40 Eth 1/1/5        14:18:77:20:86:ce  Eth 1/1/2     777   0e:fc:00:02:01:00
02:01:00 20:01:f4:e9:d4:a4:7d:c3
20:00:f4:e9:d4:a4:7d:c3
```

- To verify the name server entries on the FCF, use the `show fc ns switch brief` command.

```
FCF# show fc ns switch brief
Total number of devices = 2

Intf#                      Domain     FC-ID      Enode-WWPN              Enode-WWNN

fibrechannel1/1/3            2       02:00:00  21:04:00:11:0d:64:67:00   20:04:00:11:0d:
64:67:00
ethernet1/1/13               2       02:01:00  00:00:00:00:00:00:00:00
00:00:00:00:00:00:00:00
```

- To verify the active zoneset on the FCF, use the `show fc zoneset active` command.

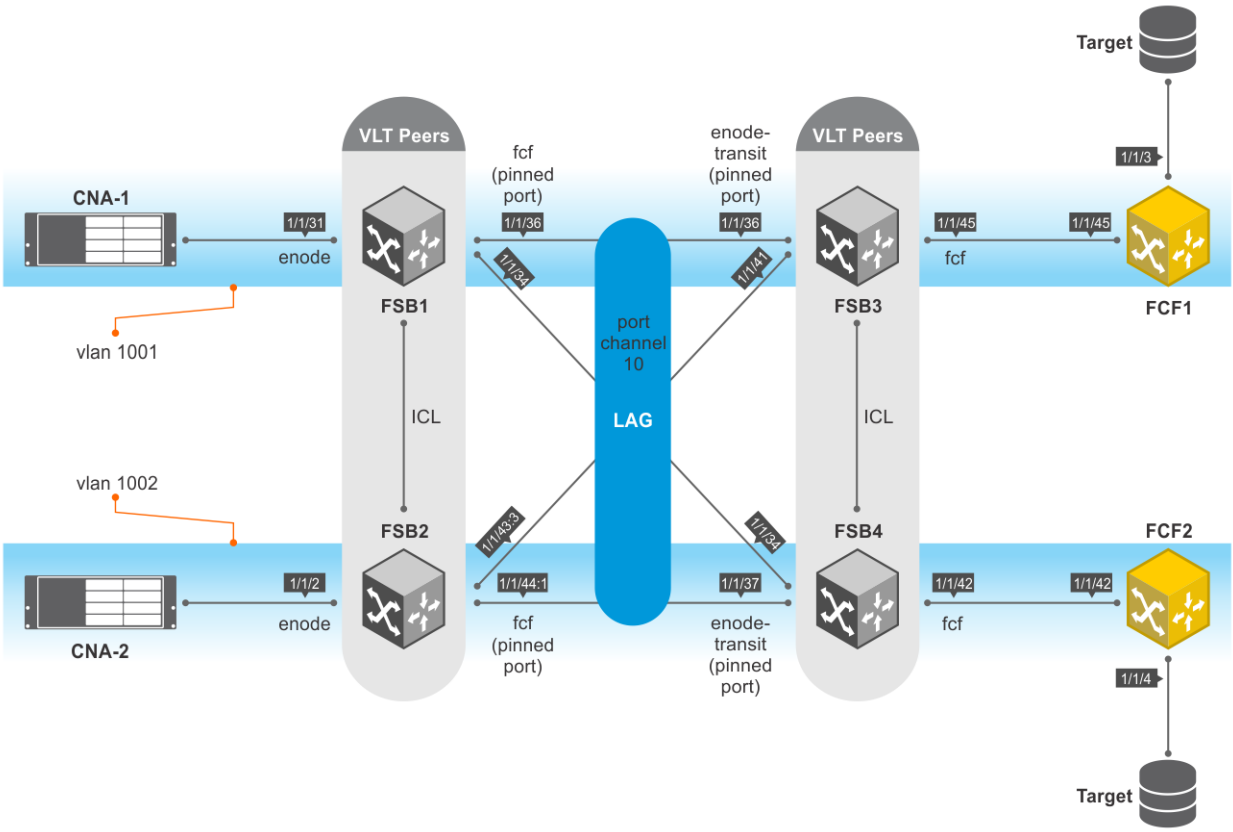
```
FCF# show fc zoneset active
```

```
vFabric id: 2
Active Zoneset: zonesetA
ZoneName                          ZoneMember
============================================================
zoneA
20:01:f4:e9:d4:a4:7d:c3

                                  21:00:00:24:ff:7c:ae:0e
```

Sample Multi-hop FSB configuration

The following is a sample multi-hop FSB topology.



In this topology:

- FSB1 and FSB2—access FSBs.
- FSB3 and FSB4—core FSBs.
- VLT is configured between FSB1 and FSB2, and requires port-pinning for VLT port channels configured between access FSBs and core FSBs. The port modes are:
  - Directly-connected CNA ports—ENode
  - Ports connected to FSB3 and FSB4—FCF
- VLT is configured between FSB3 and FSB4, and requires port-pinning for VLT port channels configured between access and core FSBs. The port modes are:
  - Ports connected to FSB1 and FSB2—ENode-transit
  - Ports connected to FCFs, for pinning to work at ENode port—FCF

The following table lists the high-level configurations on FSB1, FSB3, and FCF1. These configurations apply to FSB2, FSB4, and FCF2, respectively.

Table 3. High-level configurations on FSB1, FSB3, and FCF1

FSB1/FSB2	FSB3/FSB4	FCF1/FCF2
1 Enable FIP snooping.	1 Enable FIP snooping.	1 Enable Fiber Channel F-Port mode globally.
2 Enable DCBX.	2 Enable DCBX.	2 Create zones.
3 Create FCoE VLAN and configure FIP snooping.	3 Create FCoE VLAN and configure FIP snooping.	3 Create zoneset.
4 Create a class-map.	4 Create a class-map.	4 Create a vfabric VLAN.
5 Create a policy-map and enable PFC parameters.	5 Create a policy-map.	5 Create vfabric and activate the zoneset.
6 Configure port channel.	6 Configure port channel.	6 Enable DCBX.
7 Configure VLT domain.	7 Configure VLT domain.	7 Enable PFC parameters.
8 Configure VLTi interface member links.	8 Configure VLTi interface member links.	8 Apply PFC on the downlink interface (FSB3/FSB4).
9 Configure VLAN.	9 Configure VLAN.	9 Apply vfabric on the downlink and uplink (target-connected) interfaces.
10 Apply the policy-map parameters on uplink (FSB3/FSB4) and downlink interfaces (CNA-1/CNA-2). Configure the uplink interface as pinned-port.	10 Apply the policy-map on the uplink (FCF1/FCF2) and downlink interfaces (FSB1/FSB2). Configure the downlink interface as pinned-port.	
11 Configure FIP snooping port mode on the uplink interface.	11 Configure FIP snooping port mode on the uplink interface and the port channel.	

FSB1 configuration

1 Enable FIP snooping.

```
FSB1(config)# feature fip-snooping with-cvl
```

2 Enable DCBX.

```
FSB1(config)# dcbx enable
```

3 Create FCoE VLAN and configure FIP snooping.

```
FSB1(config)#interface vlan1001
FSB1(conf-if-vl-1001)# fip-snooping enable
FSB1(conf-if-vl-1001)# no shutdown
```

```
FSB1(config)#interface vlan1002
FSB1(conf-if-vl-1002)# fip-snooping enable
FSB1(conf-if-vl-1002)# no shutdown
```

4 Create a class-map.

```
FSB1(config)# class-map type network-qos c3
FSB1(config-cmap-nqos)# match qos-group 3
```

5 Create a policy-map.

```
FSB1(config)# policy-map type network-qos nqpolicy
FSB1(config-pmap-network-qos)# class c3
FSB1(config-pmap-c-nqos)# pause
FSB1(config-pmap-c-nqos)# pfc-cos 3
```

6 Configure port channel.

```
FSB1(config)# interface port-channel 10
FSB1(conf-if-po-10)# no shutdown
FSB1(conf-if-po-10)# vlt-port-channel 1
```

7 Configure VLT domain.

```
FSB1(config)# vlt-domain 2
FSB1(conf-vlt-2)# discovery-interface ethernet1/1/32-1/1/33
FSB1(conf-vlt-2)# vlt-mac 1a:2b:3c:0a:0b:0c
```

8 Configure VLTi interface member links.

```
FSB1(config)# interface ethernet1/1/32
FSB1(conf-if-eth1/1/32)# no shutdown
FSB1(conf-if-eth1/1/32)# no switchport
```

```
FSB1(config)# interface ethernet1/1/33
FSB1(conf-if-eth1/1/33)# no shutdown
FSB1(conf-if-eth1/1/33)# no switchport
```

```
FSB1(config)# interface ethernet1/1/34
FSB1(conf-if-eth1/1/34)# no shutdown
FSB1(conf-if-eth1/1/34)# no switchport
FSB1(conf-if-eth1/1/34)# channel-group 10
```

```
FSB1(config)# interface ethernet 1/1/36
FSB1(conf-if-eth1/1/36)# no shutdown
FSB1(conf-if-eth1/1/36)# no switchport
FSB1(conf-if-eth1/1/36)# channel-group 10
```

9 Configure VLAN on FSB1.

```
FSB1(config)# interface ethernet 1/1/31
FSB1(conf-if-eth1/1/31)# no shutdown
FSB1(conf-if-eth1/1/31)# switchport mode trunk
FSB1(conf-if-eth1/1/31)# switchport access vlan 1
FSB1(conf-if-eth1/1/31)# switchport trunk allowed vlan 1001
```

```
FSB1(config)# interface port-channel 10
FSB1(conf-if-po-10)# switchport mode trunk
FSB1(conf-if-po-10)# switchport access vlan 1
FSB1(conf-if-po-10)# switchport trunk allowed vlan 1001-1002
```

10 Apply the policy-map on the interfaces connected to FSB2 and CNA-1. Configure the interface connected to FSB2 as pinned-port.

```
FSB1(config)# interface ethernet 1/1/36
FSB1(conf-if-eth1/1/36)# flowcontrol receive off
FSB1(conf-if-eth1/1/36)# priority-flow-control mode on
FSB1(conf-if-eth1/1/36)# service-policy input type network-qos nqpolicy
FSB1(conf-if-eth1/1/36)# fcoe-pinned-port
```

```
FSB1(config)# interface ethernet 1/1/31
FSB1(conf-if-eth1/1/31)# flowcontrol receive off
FSB1(conf-if-eth1/1/31)# priority-flow-control mode on
FSB1(conf-if-eth1/1/31)# service-policy input type network-qos nqpolicy
```

11 Configure FIP snooping port mode on the port channel interface. The default port mode is ENode. Hence, the interface connected to CNA-1 does not require additional configuration.

```
FSB1(config)# interface port-channel 10
FSB1(conf-if-po-10)# fip-snooping port-mode fcf
```

FSB2 configuration

1 Enable FIP snooping.

```
FSB2(config)# feature fip-snooping with-cvl
```

2 Enable DCBX.

```
FSB2(config)# dcbx enable
```

3 Create FCoE VLAN and configure FIP snooping.

```
FSB2(config)#interface vlan1001
FSB2(conf-if-vl-1001)# fip-snooping enable
FSB2(conf-if-vl-1001)# no shutdown
```

```
FSB2(config)#interface vlan1002
FSB2(conf-if-vl-1002)# fip-snooping enable
FSB2(conf-if-vl-1002)# no shutdown
```

4 Create a class-map.

```
FSB2(config)# class-map type network-qos c3
FSB2(config-cmap-nqos)# match qos-group 3
```

5 Create a policy-map.

```
FSB2(config)# policy-map type network-qos nqpolicy
FSB2(config-pmap-network-qos)# class c3
FSB2(config-pmap-c-nqos)# pause
FSB2(config-pmap-c-nqos)# pfc-cos 3
```

6 Configure port channel.

```
FSB2(config)# interface port-channel 10
FSB2(conf-if-po-10)# no shutdown
FSB2(conf-if-po-10)# vlt-port-channel 1
```

7 Configure VLT domain.

```
FSB2(config)# vlt-domain 2
FSB2(conf-vlt-2)# discovery-interface ethernet1/1/43:1-1/1/43:2
FSB2(conf-vlt-2)# vlt-mac 1a:2b:3c:0a:0b:0c
```

8 Configure VLTi interface member links.

```
FSB2(config)# interface ethernet1/1/43:1
FSB2(conf-if-eth1/1/43:1)# no shutdown
FSB2(conf-if-eth1/1/43:1)# no switchport
```

```
FSB2(config)# interface ethernet1/1/43:2
FSB2(conf-if-eth1/1/43:2)# no shutdown
FSB2(conf-if-eth1/1/43:2)# no switchport
```

```
FSB2(config)# interface ethernet 1/1/43:3
FSB2(conf-if-eth1/1/43:3)# no shutdown
FSB2(conf-if-eth1/1/43:3)# no switchport
FSB2(conf-if-eth1/1/43:3)# channel-group 10
```

```
FSB2(config)# interface ethernet1/1/44:1
FSB2(conf-if-eth1/1/44:1)# no shutdown
FSB2(conf-if-eth1/1/44:1)# no switchport
FSB2(conf-if-eth1/1/44:1)# channel-group 10
```

9 Configure VLAN on FSB2.

```
FSB2(config)# interface ethernet 1/1/2
FSB2(conf-if-eth1/1/2)# no shutdown
FSB2(conf-if-eth1/1/2)# switchport mode trunk
FSB2(conf-if-eth1/1/2)# switchport access vlan 1
FSB2(conf-if-eth1/1/2)# switchport trunk allowed vlan 1002
```

```
FSB2(config)# interface port-channel 10
FSB2(conf-if-po-10)# switchport mode trunk
FSB2(conf-if-po-10)# switchport access vlan 1
FSB2(conf-if-po-10)# switchport trunk allowed vlan 1001-1002
```

10 Apply the policy-map on the interfaces connected to FSB4 and CNA-2. Configure the interface connected to FSB4 as pinned-port.

```
FSB2(config)# interface ethernet 1/1/44:1
FSB2(conf-if-eth1/1/44:1)# flowcontrol receive off
FSB2(conf-if-eth1/1/44:1)# priority-flow-control mode on
FSB2(conf-if-eth1/1/44:1)# service-policy input type network-qos nqpolicy
FSB2(conf-if-eth1/1/44:1)# fcoe-pinned-port
```

```
FSB2(config)# interface ethernet 1/1/2
FSB2(conf-if-eth1/1/2)# flowcontrol receive off
FSB2(conf-if-eth1/1/2)# priority-flow-control mode on
FSB2(conf-if-eth1/1/2)# service-policy input type network-qos nqpolicy
```

11 Configure FIP snooping port mode on the port channel interface. The default port mode is ENode. Hence, the interface connected to CNA-2 does not require additional configuration.

```
FSB2(config)# interface port-channel 10
FSB2(conf-if-po-10)# fip-snooping port-mode fcf
```

FSB3 configuration

1 Enable FIP snooping.

```
FSB3(config)# feature fip-snooping with-cvl
```

2 Enable DCBX.

```
FSB3(config)# dcbx enable
```

3 Create FCoE VLAN and configure FIP snooping.

```
FSB3(config)#interface vlan1001
FSB3(conf-if-vl-1001)# fip-snooping enable
FSB3(conf-if-vl-1001)# no shutdown
```

```
FSB3(config)#interface vlan1002
FSB3(conf-if-vl-1002)# fip-snooping enable
FSB3(conf-if-vl-1002)# no shutdown
```

4 Create a class-map.

```
FSB3(config)# class-map type network-qos c3
FSB3(config-cmap-nqos)# match qos-group 3
```

5 Create a policy-map.

```
FSB3(config)# policy-map type network-qos nqpolicy
FSB3(config-pmap-network-qos)# class c3
FSB3(config-pmap-c-nqos)# pause
FSB3(config-pmap-c-nqos)# pfc-cos 3
```

6 Configure port channel.

```
FSB3(config)# interface port-channel 10
FSB3(conf-if-po-10)# no shutdown
FSB3(conf-if-po-10)# vlt-port-channel 1
```

7 Configure VLT domain.

```
FSB3(config)# vlt-domain 3
FSB3(conf-vlt-3)# discovery-interface ethernet1/1/39-1/1/40
FSB3(conf-vlt-3)# vlt-mac 1a:2b:3c:2a:1b:1c
```

8 Configure VLTi interface member links.

```
FSB3(config)# interface ethernet1/1/39
FSB3(conf-if-eth1/1/39)# no shutdown
FSB3(conf-if-eth1/1/39)# no switchport
```

```
FSB3(config)# interface ethernet1/1/40
FSB3(conf-if-eth1/1/40)# no shutdown
FSB3(conf-if-eth1/1/40)# no switchport
```

```
FSB3(config)# interface ethernet1/1/41
FSB3(conf-if-eth1/1/41)# no shutdown
FSB3(conf-if-eth1/1/41)# no switchport
FSB3(conf-if-eth1/1/41)# channel-group 10
```

```
FSB3(config)# interface ethernet 1/1/36
FSB3(conf-if-eth1/1/36)# no shutdown
FSB3(conf-if-eth1/1/36)# no switchport
FSB3(conf-if-eth1/1/36)# channel-group 10
```

9 Configure VLAN on FSB3.

```
FSB3(config)# interface ethernet 1/1/45
FSB3(conf-if-eth1/1/45)# no shutdown
FSB3(conf-if-eth1/1/45)# switchport mode trunk
FSB3(conf-if-eth1/1/45)# switchport access vlan 1
FSB3(conf-if-eth1/1/45)# switchport trunk allowed vlan 1001
```

```
FSB3(config)# interface port-channel 10
FSB3(conf-if-po-10)# switchport mode trunk
FSB3(conf-if-po-10)# switchport access vlan 1
FSB3(conf-if-po-10)# switchport trunk allowed vlan 1001-1002
```

10 Apply the policy-map on the interfaces connected to FCF1 and FSB1. Configure the interface connected to FSB1 as pinned-port.

```
FSB3(config)# interface ethernet 1/1/45
FSB3(conf-if-eth1/1/45)# flowcontrol receive off
FSB3(conf-if-eth1/1/45)# priority-flow-control mode on
FSB3(conf-if-eth1/1/45)# service-policy input type network-qos nqpolicy
```

```
FSB3(config)# interface ethernet 1/1/36
FSB3(conf-if-eth1/1/36)# flowcontrol receive off
```

```
FSB3(conf-if-eth1/1/36)# priority-flow-control mode on
FSB3(conf-if-eth1/1/36)# service-policy input type network-qos nqpolicy
FSB3(conf-if-eth1/1/36)# fcoe-pinned-port
```

11 Configure FIP snooping port mode on the port channel and the interface connected to FCF1.

```
FSB3(config)# interface port-channel 10
FSB3(conf-if-po-10)# fip-snooping port-mode enode-transit
```

```
FSB3(config)# interface ethernet 1/1/45
FSB3(conf-if-eth1/1/45)# fip-snooping port-mode fcf
```

FSB4 configuration

1 Enable FIP snooping.

```
FSB4(config)# feature fip-snooping with-cvl
```

2 Enable DCBX.

```
FSB4(config)# dcbx enable
```

3 Create FCoE VLAN and configure FIP snooping.

```
FSB4(config)#interface vlan1001
FSB4(conf-if-vl-1001)# fip-snooping enable
FSB4(conf-if-vl-1001)# no shutdown
```

```
FSB4(config)#interface vlan1002
FSB4(conf-if-vl-1002)# fip-snooping enable
FSB4(conf-if-vl-1002)# no shutdown
```

4 Create a class-map.

```
FSB4(config)# class-map type network-qos c3
FSB4(config-cmap-nqos)# match qos-group 3
```

5 Create a policy-map.

```
FSB4(config)# policy-map type network-qos nqpolicy
FSB4(config-pmap-network-qos)# class c3
FSB4(config-pmap-c-nqos)# pause
FSB4(config-pmap-c-nqos)# pfc-cos 3
```

6 Configure port channel.

```
FSB4(config)# interface port-channel 10
FSB4(conf-if-po-10)# no shutdown
FSB4(conf-if-po-10)# vlt-port-channel 1
```

7 Configure VLT domain.

```
FSB4(config)# vlt-domain 3
FSB4(conf-vlt-2)# discovery-interface ethernet1/1/40
FSB4(conf-vlt-2)# vlt-mac 1a:2b:3c:2a:1b:1c
```

8 Configure VLTi interface member links.

```
FSB4(config)# interface ethernet1/1/34
FSB4(conf-if-eth1/1/34)# no shutdown
FSB4(conf-if-eth1/1/34)# no switchport
FSB4(conf-if-eth1/1/34)# channel-group 10
```

```
FSB4(config)# interface ethernet1/1/37
FSB4(conf-if-eth1/1/37)# no shutdown
FSB4(conf-if-eth1/1/37)# no switchport
FSB4(conf-if-eth1/1/37)# channel-group 10
```

9 Configure VLAN on FSB4.

```
FSB4(config)# interface ethernet 1/1/42
FSB4(conf-if-eth1/1/42)# no shutdown
FSB4(conf-if-eth1/1/42)# switchport mode trunk
FSB4(conf-if-eth1/1/42)# switchport access vlan 1
FSB4(conf-if-eth1/1/42)# switchport trunk allowed vlan 1002
```

```
FSB4(config)# interface port-channel 10
FSB4(conf-if-po-10)# switchport mode trunk
```

```
FSB4(conf-if-po-10)# switchport access vlan 1
FSB4(conf-if-po-10)# switchport trunk allowed vlan 1001-1002
```

10 Apply the policy-map on the interfaces connected to FCF2.

```
FSB4(config)# interface ethernet 1/1/42
FSB4(conf-if-eth1/1/42)# flowcontrol receive off
FSB4(conf-if-eth1/1/42)# priority-flow-control mode on
FSB4(conf-if-eth1/1/42)# service-policy input type network-qos nqpolicy
```

11 Configure FIP snooping port mode on the port channel and the interface connected to FCF2. Configure the interface connected to FSB2 as pinned-port.

```
FSB4(config)# interface port-channel 10
FSB4(conf-if-po-10)# fip-snooping port-mode enode-transit
```

```
FSB4(config)# interface ethernet 1/1/42
FSB4(conf-if-eth1/1/42)# fip-snooping port-mode fcf
```

```
FSB4(config)# interface ethernet 1/1/37
FSB4(conf-if-eth1/1/37)# fcoe-pinned-port
```

FCF1 configuration

1 Enable Fiber Channel F-Port mode globally.

```
FCF1(config)# feature fc domain-id 2
```

2 Create zones.

```
FCF1(config)# fc zone zoneA
FCF1(config-fc-zone-zoneA)# member wwn 23:05:22:11:0d:64:67:11
FCF1(config-fc-zone-zoneA)# member wwn 22:04:22:13:0d:64:67:00
```

3 Create zoneset.

```
FCF1(config)# fc zoneset zonesetA
FCF1(conf-fc-zoneset-setA)# member zoneA
```

4 Create a vfabric VLAN.

```
FCF1(config)# interface vlan 1001
```

5 Create vfabric and activate the zoneset.

```
FCF1(config)# vfabric 1
FCF1(conf-vfabric-1)# vlan 1001
FCF1(conf-vfabric-1)# fcoe fcmap 0xEFC00
FCF1(conf-vfabric-1)# zoneset activate zonesetA
```

6 Enable DCBX.

```
FCF1(config)# dcbx enable
```

7 Enable PFC parameters.

```
FCF1(config)# class-map type network-qos c3
FCF1(config-cmap-nqos)# match qos-group 3
FCF1(config-cmap-nqos)# exit
FCF1(config)# policy-map type network-qos nqpolicy
FCF1(config-pmap-network-qos)# class c3
FCF1(config-pmap-c-nqos)# pause
FCF1(config-pmap-c-nqos)# pfc-cos 3
```

8 Apply PFC on the interface connected to FSB3.

```
FCF1(config)# interface ethernet 1/1/45
FCF1(conf-if-eth1/1/45)# no shutdown
FCF1(conf-if-eth1/1/45)# flowcontrol receive off
FCF1(conf-if-eth1/1/45)# priority-flow-control mode on
FCF1(conf-if-eth1/1/45)# service-policy input type network-qos nqpolicy
```

9 Apply vfabric on the interfaces connected to FSB3 and the target.

```
FCF1(config)# interface ethernet 1/1/45
FCF1(conf-if-eth1/1/45)# switchport access vlan 1
FCF1(conf-if-eth1/1/45)# vfabric 1
```

```
FCF1(config)# interface fibrechannel 1/1/3
FCF1(conf-if-fc1/1/3)# description target_connected_port
```

```
FCF1(conf-if-fc1/1/3)# no shutdown
FCF1(conf-if-fc1/1/3)# vfabric 1
```

FCF2 configuration

1 Enable Fiber Channel F-Port mode globally.

```
FCF2(config)# feature fc domain-id 3
```

2 Create zones.

```
FCF2(config)# fc zone zoneB
FCF2(config-fc-zone-zoneB)# member wwn 20:01:00:0e:1e:f1:f1:84
FCF2(config-fc-zone-zoneB)# member wwn 20:00:00:0e:1e:f1:f1:84
```

3 Create zoneset.

```
FCF2(config)# fc zoneset zonesetB
FCF2(conf-fc-zoneset-setB)# member zoneB
```

4 Create a vfabric VLAN.

```
FCF2(config)# interface vlan 1002
```

5 Create vfabric and activate the zoneset.

```
FCF2(config)# vfabric 2
FCF2(conf-vfabric-2)# vlan 1002
FCF2(conf-vfabric-2)# fcoe fcmap 0xEFC00
FCF2(conf-vfabric-2)# zoneset activate zonesetB
```

6 Enable DCBX.

```
FCF2(config)# dcbx enable
```

7 Enable PFC parameters.

```
FCF2(config)# class-map type network-qos c3
FCF2(config-cmap-nqos)# match qos-group 3
FCF2(config-cmap-nqos)# exit
FCF2(config)# policy-map type network-qos nqpolicy
FCF2(config-pmap-network-qos)# class c3
FCF2(config-pmap-c-nqos)# pause
FCF2(config-pmap-c-nqos)# pfc-cos 3
```

8 Apply PFC on the interface connected to FSB4.

```
FCF2(config)# interface ethernet 1/1/42
FCF2(conf-if-eth1/1/42)# no shutdown
FCF2(conf-if-eth1/1/42)# flowcontrol receive off
FCF2(conf-if-eth1/1/42)# priority-flow-control mode on
FCF2(conf-if-eth1/1/42)# service-policy input type network-qos nqpolicy
```

9 Apply vfabric on the interfaces connected to FSB4 and the target.

```
FCF2(config)# interface ethernet 1/1/42
FCF2(conf-if-eth1/1/42)# switchport access vlan 1
FCF2(conf-if-eth1/1/42)# vfabric 1
```

```
FCF2(config)# interface fibrechannel 1/1/4
FCF2(conf-if-fc1/1/4)# description target_connected_port
FCF2(conf-if-fc1/1/4)# no shutdown
FCF2(conf-if-fc1/1/4)# vfabric 2
```

Verify the configuration

Use the following `show` commands to verify the configuration:

FSB1

```
FSB1# show fcoe sessions
Enode MAC           Enode Interface FCF MAC          FCF interface      VLAN   FCoE
MAC          FC-ID   PORT WWPN
PORT WWNN
```

```
--------------------------------------------------------------------------------------------------------
------------------------------
f4:e9:d4:f9:fc:42   Eth 1/1/31      14:18:77:20:86:ce Po 10(Eth 1/1/36) 1001 0e:fc:00:02:02:00
02:02:00 23:05:22:11:0d:64:67:11   22:04:22:13:0d:64:67:00
```

```
FSB1# show fcoe fcf
FCF MAC             FCF Interface       VLAN     FC-MAP         FKA_ADV_PERIOD    No. of
Enodes      FCF Mode
--------------------------------------------------------------------------------------------------------
--------------
14:18:77:20:86:ce   Po 10(Eth 1/1/36)    1001     0e:fc:00       8000
1                   F
```

```
FSB1# show fcoe system
Mode                          : FSB
CVL Status                    : Enabled
FCOE VLAN List (Operational) : 1001,1002
FCFs                          : 1
Enodes                        : 1
Sessions                      : 1
```

FSB2

```
FSB2# show fcoe sessions
Enode MAC         Enode Interface FCF MAC          FCF interface      VLAN FCoE MAC          FC-
ID    PORT WWPN                PORT WWNN
--------------------------------------------------------------------------------------------------------
--------------------------------------------------
00:0e:1e:f1:f1:84   Eth 1/1/1   14:18:77:20:80:ce  Po 10(Eth 1/1/44:1)1002 0e:fc:00:02:01:00
02:01:00 20:01:00:0e:1e:f1:f1:84  20:00:00:0e:1e:f1:f1:84
```

```
FSB2# show fcoe fcf
FCF MAC             FCF Interface       VLAN     FC-MAP         FKA_ADV_PERIOD    No. of
Enodes      FCF Mode
--------------------------------------------------------------------------------------------------------
--------------
14:18:77:20:80:ce   Po 10(Eth 1/1/44:1)  1002     0e:fc:00       8000
1                   F
```

```
FSB2# show fcoe system
Mode                          : FSB
CVL Status                    : Enabled
FCOE VLAN List (Operational) : 1001,1002
FCFs                          : 1
Enodes                        : 1
Sessions                      : 1
```

FSB3

```
FSB3# show fcoe sessions
Enode MAC         Enode Interface   FCF MAC           FCF interface VLAN FCoE MAC          FC-
ID    PORT WWPN
PORT WWNN
--------------------------------------------------------------------------------------------------------
---------------------------
f4:e9:d4:f9:fc:42 Po 10(Eth 1/1/36) 14:18:77:20:86:ce Eth 1/1/45    1001 0e:fc:00:02:02:00
02:02:00 23:05:22:11:0d:64:67:11   22:04:22:13:0d:64:67:00
```

```
FSB3# show fcoe fcf
FCF MAC             FCF Interface       VLAN     FC-MAP         FKA_ADV_PERIOD    No. of
Enodes      FCF Mode
--------------------------------------------------------------------------------------------------------
--------------
14:18:77:20:86:ce   Eth 1/1/45           1001     0e:fc:00       8000
1                   F
```

```
FSB3# show fcoe system
Mode                          : FSB
CVL Status                    : Enabled
FCOE VLAN List (Operational) : 1001,1002
```

```
FCFs                          : 1
Enodes                        : 1
Sessions                      : 1
```

FSB4

```
FSB4# show fcoe sessions
Enode MAC           Enode Interface      FCF MAC           FCF interface      VLAN      FCoE
MAC          FC-ID   PORT WWPN               PORT WWNN
------------------------------------------------------------------------------------------
-----------------------------------------------
00:0e:1e:f1:f1:84   Po 10(Eth 1/1/37)    14:18:77:20:80:ce  Eth 1/1/42           1002   0e:fc:
00:02:01:00  02:01:00 20:01:00:0e:1e:f1:f1:84  20:00:00:0e:1e:f1:f1:84
```

```
FSB4# show fcoe fcf
FCF MAC             FCF Interface       VLAN     FC-MAP         FKA_ADV_PERIOD    No. of
Enodes      FCF Mode
------------------------------------------------------------------------------------------
--------------
14:18:77:20:80:ce   Eth 1/1/42           1002     0e:fc:00       8000
1                   F
```

```
FSB4# show fcoe system
Mode                          : FSB
CVL Status                    : Enabled
FCOE VLAN List (Operational) : 1001,1002
FCFs                          : 1
Enodes                        : 1
Sessions                      : 1
```

FCF1

```
FCF1# show fcoe sessions
Enode MAC          Enode Interface FCF MAC           FCF interface VLAN   FCoE MAC          FC-
ID    PORT WWPN               PORT WWNN
------------------------------------------------------------------------------------------
-----------------------------------------------
f4:e9:d4:f9:fc:42 Eth 1/1/45     14:18:77:20:86:ce  ~             1001 0e:fc:00:02:02:00
02:02:00 23:05:22:11:0d:64:67:11 22:04:22:13:0d:64:67:00
```

```
FCF1# show fc ns switch brief
Total number of devices = 2

Intf#                     Domain     FC-ID      Enode-WWPN              Enode-WWNN

fibrechannel1/1/3           2       02:00:00  21:04:00:11:0d:64:67:00   20:04:00:11:0d:64:67:00
ethernet1/1/45              2       02:02:00  23:05:22:11:0d:64:67:11   22:04:22:13:0d:64:67:00
```

FCF2

```
FCF2# show fcoe sessions
Enode MAC          Enode Interface FCF MAC       FCF interface VLAN      FCoE MAC          FC-
ID   PORT WWPN              PORT WWNN
------------------------------------------------------------------------------------------
-----------------------------------------------
00:0e:1e:f1:f1:84 Eth 1/1/42    14:18:77:20:80:ce  ~          1002   0e:fc:00:02:01:00
02:00:01 20:01:00:0e:1e:f1:f1:84  20:00:00:0e:1e:f1:f1:84
```

```
FCF2# show fc ns switch brief
Total number of devices = 1

Intf#                     Domain     FC-ID      Enode-WWPN              Enode-WWNN

fibrechannel1/1/4           3       02:01:00  22:04:00:11:0d:64:67:00   22:04:00:11:0d:64:67:00
ethernet1/1/42              3       02:00:01  20:01:00:0e:1e:f1:f1:84  20:00:00:0e:1e:f1:f1:84
```

Configuration guidelines

When configuring different modes; for example, F_Port, NPG, or FSB, consider the following:

- F_Port, NPG, and FSB modes are mutually exclusive. You can enable only one at a time.
- You can enable the mode-specific commands only after enabling the specific feature.
- Before you disable the F_Port and NPG features, delete the mode-specific configurations. When you disable FSB, the system automatically removes the configurations.

F_Port commands

The following commands are supported on F_Port mode:

fc alias

Creates an FC alias. After creating the alias, add members to the FC alias. An FC alias can have a maximum of 255 unique members.

Syntax	`fc alias` *`alias-name`*
Parameters	*`alias-name`* — Enter a name for the FC alias.
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command deletes the FC alias. To delete an FC alias, first remove it from the FC zone.

Example	`OS10(config)# fc alias test`
	`OS10(config-fc-alias-test)# member wwn 21:00:00:24:ff:7b:f5:c9`
	`OS10(config-fc-alias-test)# member wwn 20:25:78:2b:cb:6f:65:57`

Supported Releases	10.3.1E or later

fc zone

Creates an FC zone and adds members to the zone. An FC zone can have a maximum of 255 unique members.

Syntax	`fc zone` *`zone-name`*
Parameters	*`zone-name`* — Enter a name for the zone.
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command deletes the FC zone. To delete an FC zone, first remove it from the FC zoneset.

Example	`OS10(config)# fc zone hba1`
	`OS10(config-fc-zone-hba1)# member wwn 10:00:00:90:fa:b8:22:19`
	`OS10(config-fc-zone-hba1)# member wwn 21:00:00:24:ff:7b:f5:c8`

Supported Releases	10.3.1E or later

fc zoneset

Creates an FC zoneset and adds the existing FC zones to the zoneset.

Syntax `fc zoneset zoneset-name`

Parameters *zoneset-name* — Enter a name for the FC zoneset. The name must start with a letter and may contain these characters: A-Z, a-z, 0-9, $, _, -, ^

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the FC zoneset.

Example

```
OS10(config)# fc zoneset set
OS10(conf-fc-zoneset-set)# member hba1
```

Supported Releases 10.3.1E or later

feature fc

Enables the F_Port globally.

Syntax `feature fc domain-id domain-id`

Parameters *domain-id* — Enter the domain ID of the F_Port, from 1 to 239.

Defaults Disabled

Command Mode CONFIGURATION

Usage Information The `no` version of this command disables the F_Port. You can disable the F_Port only when vfabric and zoning configurations are not available. Before disabling the F_Port, remove the vfabric and zoning configurations. You can enable only one of the following at a time: F_Port, NPG, or FSB.

Example

```
OS10(config)# feature fc domain-id 100
```

Supported Releases 10.3.1E or later

member (alias)

Add members to existing FC aliases. Identify a member by an FC alias, a world wide name (WWN), or an FC ID.

Syntax `member {wwn wwn-ID | fc-id fc-id}`

Parameters

- *wwn-ID* — Enter the WWN name.
- *fc-id* — Enter the FC ID name.

Defaults Not configured

Command Mode Alias CONFIGURATION

Usage Information The `no` version of this command removes the member from the FC alias.

Example

```
OS10(config)# fc alias test
OS10(config-fc-alias-test)# member wwn 21:00:00:24:ff:7b:f5:c9
OS10(config-fc-alias-test)# member wwn 20:25:78:2b:cb:6f:65:57
```

Supported Releases 10.3.1E or later

member (zone)

Adds members to existing zones. Identify a member by an FC alias, a world wide name (WWN), or an FC ID.

Syntax `member {alias-name alias-name | wwn wwn-ID | fc-id fc-id}`

Parameters

- *alias-name* — Enter the FC alias name.
- *wwn-ID* — Enter the WWN name.
- *fc-id* — Enter the FC ID name.

Defaults Not configured

Command Mode Zone CONFIGURATION

Usage Information The `no` version of this command removes the member from the zone.

Example

```
OS10(config)# fc zone hba1
OS10(config-fc-zone-hba1)# member wwn 10:00:00:90:fa:b8:22:19
OS10(config-fc-zone-hba1)# member wwn 21:00:00:24:ff:7b:f5:c8
```

Supported Releases 10.3.1E or later

member (zoneset)

Adds zones to an existing zoneset.

Syntax `member zone-name`

Parameters *zone-name* — Enter an existing zone name.

Defaults Not configured

Command Mode Zoneset CONFIGURATION

Usage Information The `no` version of this command removes the zone from the zoneset.

Example

```
OS10(config)# fc zoneset set
OS10(conf-fc-zoneset-set)# member hba1
```

Supported Releases 10.3.1E or later

show fc alias

Displays the details of a FC alias and its members.

Syntax `show fc alias [alias-name]`

Parameters *alias-name* — (Optional) Enter the FC alias name.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show fc alias

Alias Name               Alias Member
================================================
test                     21:00:00:24:ff:7b:f5:c9
                         20:25:78:2b:cb:6f:65:57

OS10#
```

Supported Releases 10.3.1E or later

show fc interface-area-id mapping

Displays the FC ID to interface mapping details.

Syntax `show fc interface-area-id mapping`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show fc interface-area-id mapping
Intf Name                 FC-ID          Status
================================================
ethernet1/1/40           0a:02:00        Active
```

Supported Releases 10.4.1.0 or later

show fc ns switch

Displays the details of the FC NS switch parameters.

Syntax `show fc ns switch [brief]`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show fc ns switch

Total number of devices = 1
Switch Name                 10:00:14:18:77:13:38:28
Domain Id                   4
Switch Port                 port-channel10(Eth 1/1/9)
FC-Id                       04:00:00
Port Name                   50:00:d3:10:00:ec:f9:05
Node Name                   50:00:d3:10:00:ec:f9:00
Class of Service            8
Symbolic Port Name          Compellent Port QLGC FC 8Gbps; Slot=06 Port=01 in
Controller: SN 60665 of Storage Center: DEVTEST 60665
Symbolic Node Name          Compellent Storage Center: DEVTEST 60665
Port Type                   N_PORT
Registered with NameServer  Yes
Registered for SCN          No
```

Example (brief)

```
OS10# show fc ns switch brief
Total number of devices = 1
Intf#                       Domain    FC-ID      Enode-WWPN              Enode-
WWNN
port-channel10(Eth 1/1/9)      4     04:00:00   10:00:00:90:fa:b8:22:18
20:00:00:90:fa:b8:22:18
```

Supported Releases 10.3.1E or later

show fc zone

Displays the FC zones and the zone members.

Syntax `show fc zone [zone-name]`

Parameters *zone-name* — Enter the FC zone name.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show fc zone

Zone Name                   Zone Member
================================================
hba1                       21:00:00:24:ff:7b:f5:c8
                           10:00:00:90:fa:b8:22:19
                           21:00:00:24:ff:7f:ce:ee
                           21:00:00:24:ff:7f:ce:ef
hba2                       20:01:00:0e:1e:e8:e4:99
                           50:00:d3:10:00:ec:f9:1b
                           50:00:d3:10:00:ec:f9:05
                           50:00:d3:10:00:ec:f9:1f
                           20:35:78:2b:cb:6f:65:57
```

Example (with zone name)

```
OS10# show fc zone hba1

Zone Name                   Zone Member
================================================
hba1                       21:00:00:24:ff:7b:f5:c8
                           10:00:00:90:fa:b8:22:19
```

```
                                21:00:00:24:ff:7f:ce:ee
                                21:00:00:24:ff:7f:ce:ef
```

Supported Releases 10.3.1E or later

show fc zoneset

Displays the FC zonesets, the zones in the zoneset, and the zone members.

Syntax `show fc zoneset [active | zoneset-name]`

Parameters *zoneset-name* — Enter the FC zoneset name.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show fc zoneset
ZoneSetName      ZoneName          ZoneMember
=========================================================
set              hba1              21:00:00:24:ff:7b:f5:c8
                                   10:00:00:90:fa:b8:22:19
                                   21:00:00:24:ff:7f:ce:ee
                                   21:00:00:24:ff:7f:ce:ef

                 hba2              20:01:00:0e:1e:e8:e4:99
                                   50:00:d3:10:00:ec:f9:1b
                                   50:00:d3:10:00:ec:f9:05
                                   50:00:d3:10:00:ec:f9:1f
                                   20:35:78:2b:cb:6f:65:57

vFabric id: 100
Active Zoneset: set
ZoneName                ZoneMember
=============================================
hba2                    20:01:00:0e:1e:e8:e4:99
                        20:35:78:2b:cb:6f:65:57
                        50:00:d3:10:00:ec:f9:05
                        50:00:d3:10:00:ec:f9:1b
                        50:00:d3:10:00:ec:f9:1f

hba1                   *10:00:00:90:fa:b8:22:19
                       *21:00:00:24:ff:7b:f5:c8
                        21:00:00:24:ff:7f:ce:ee
                        21:00:00:24:ff:7f:ce:ef
```

Example (active zoneset)

```
OS10# show fc zoneset active

vFabric id: 100
Active Zoneset: set
ZoneName                          ZoneMember
=============================================================
hba2                              20:01:00:0e:1e:e8:e4:99
                                  20:35:78:2b:cb:6f:65:57
                                  50:00:d3:10:00:ec:f9:05
                                  50:00:d3:10:00:ec:f9:1b
                                  50:00:d3:10:00:ec:f9:1f

hba1                             *10:00:00:90:fa:b8:22:19
                                 *21:00:00:24:ff:7b:f5:c8
```

```
                                    21:00:00:24:ff:7f:ce:ee
                                    21:00:00:24:ff:7f:ce:ef
```

Example (with zoneset name)

```
OS10# show fc zoneset set
ZoneSetName             ZoneName             ZoneMember
=================================================================
set                     hba1                 21:00:00:24:ff:7b:f5:c8
                                             10:00:00:90:fa:b8:22:19
                                             21:00:00:24:ff:7f:ce:ee
                                             21:00:00:24:ff:7f:ce:ef

                        hba2                 20:01:00:0e:1e:e8:e4:99
                                             50:00:d3:10:00:ec:f9:1b
                                             50:00:d3:10:00:ec:f9:05
                                             50:00:d3:10:00:ec:f9:1f
                                             20:35:78:2b:cb:6f:65:57
```

Supported Releases 10.3.1E or later

zone default-zone permit

Enables access between all logged-in FC nodes of the vfabric in the absence of an active zoneset configuration.

Syntax `zone default-zone permit`

Parameters None

Defaults Not configured

Command Mode Vfabric CONFIGURATION

Usage Information A default zone advertises a maximum of 255 members in the registered state change notification (RSCN) message. The `no` version of this command disables access between the FC nodes in the absence of an active zoneset.

Example

```
OS10(config)# vfabric 100
OS10(conf-vfabric-100)# zone default-zone permit
```

Supported Releases 10.3.1E or later

zoneset activate

Activates an existing zoneset. You can activate only one zoneset in a vfabric.

Syntax `zoneset activate` *`zoneset-name`*

Parameters *`zoneset-name`* — Enter an existing zoneset name.

Defaults Not configured

Command Mode Vfabric CONFIGURATION

Usage Information After you disable an active zoneset, the `zone default-zone permit` command configuration takes effect. Based on this configuration, the default zone allows or denies access between all the logged-in FC nodes of the vfabric. The `no` version of this command deactivates the zoneset.

Example	`OS10(config)# vfabric 100` `OS10(conf-vfabric-100)# zoneset activate set`
Supported Releases	10.3.1E or later

NPG commands

The following commands are supported on NPG mode:

fc port-mode F

Configures port mode on Fibre Channel interfaces.

Syntax	`fc port-mode F`
Parameters	None
Defaults	N_Port
Command Mode	Fibre Channel INTERFACE
Usage Information	Configure the port mode when the port is in Shut mode and when NPG mode is enabled. The `no` version of this command returns the port mode to default.
Example	`OS10(config)# interface fibrechannel 1/1/1` `OS10(conf-if-fc1/1/1)# fc port-mode F`
Supported Releases	10.4.1.0 or later

feature fc npg

Enables the NPG mode globally.

Syntax	`feature fc npg`
Parameters	None
Defaults	Disabled
Command Mode	CONFIGURATION
Usage Information	You can enable only one of the following at a time: F_Port, NPG, or FSB. The `no` version of this command disables NPG mode.
Example	`OS10(config)# feature fc npg`
Supported Releases	10.4.0E(R1) or later

show npg devices

Displays the NPG devices connected to the switch.

Syntax	`show npg devices [brief]`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information Use the `brief` option to display minimum details.

Example

```
OS10# show npg devices
Enode[0]:
Enode MAC :d4:ae:52:1a:ee:54
Enode Interface: port-channel10(Eth 1/1/9)
FCF MAC :14:18:77:20:7c:e3
Fabric Interface :Fc 1/1/25
FCoE Vlan :1001
Vfabric Id :10
ENode WWPN :20:01:d4:ae:52:1a:ee:54
ENode WWNN :20:00:d4:ae:52:1a:ee:54
FCoE MAC :0e:fc:00:01:04:02
FC-ID :01:04:02
Login Method :FLOGI
Time since discovered(in Secs) :6253
Status :LOGGED_IN
```

Example (brief)

```
Total NPG Devices = 1
ENode-Interface   ENode-WWPN              FCoE-Vlan  Fabric-Intf   Vfabric-Id  LoginMet
-------------------------------------------------------------------------------------
Po 10(Eth 1/1/9)  20:01:d4:ae:52:1a:ee:54  1001      Fc 1/1/25     10          FLOGI
LOGGED_IN
```

Supported Releases 10.4.0E(R1) or later

F_Port and NPG commands

The following commands are supported on both F_Port and NPG modes:

clear fc statistics

Clears FC statistics for specified vfabric or fibre channel interface.

Syntax `clear fc statistics [vfabric` *`vfabric-ID`* `| interface` *`fibrechannel`*`]`

Parameters

- *`vfabric-ID`* — Enter the vfabric ID.
- *`fibrechannel`* — Enter the fibre channel interface name.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# clear fc statistics vfabric 100
OS10# clear fc statistics interface fibrechannel1/1/25
```

Supported Releases 10.4.1.0 or later

fcoe

Adds FCoE parameters to the vfabric.

Syntax

```
fcoe {fcmap fc-map | fcf-priority fcf-priority-value | fka-adv-period adv-period | vlan-priority vlan-priority-value | keep-alive}
```

Parameters

- *fc-map* — Enter the FC map ID, from `0xefc00` to `0xefcff`.
- *fcf-priority-value* — Enter the FCF priority value, from 1 to 255.
- *adv-period* — Enter the FCF keepalive advertisement period, from 8 to 90 seconds.
- *vlan-priority-value* — Enter the VLAN priority value, from 0 to 7.

Defaults

- `fcmap`—0x0EFC00
- `fcf-priority`—128
- `fka-adv-period`—8
- `vlan-priority`—3
- `keep-alive`—True

Command Mode Vfabric CONFIGURATION

Usage Information The `no` version of this command disables the FCoE parameters.

Example

```
OS10(config)# vfabric 10
OS10(conf-vfabric-10)# name 10
OS10(conf-vfabric-10)# fcoe fcmap 0x0efc01
OS10(conf-vfabric-10)# fcoe fcf-priority 128
OS10(conf-vfabric-10)# fcoe fka-adv-period 8
OS10(conf-vfabric-10)# fcoe vlan-priority 3
```

Supported Releases 10.3.1E or later

name

Configures a vfabric name.

Syntax

```
name vfabric-name
```

Parameters *vfabric-name* — Enter a name for the vfabric.

Defaults Not configured

Command Mode Vfabric CONFIGURATION

Usage Information The `no` version of this command removes the vfabric name..

Example

```
OS10(config)# vfabric 100
OS10(conf-vfabric-100)# name test_vfab
```

Supported Releases 10.3.1E or later

show fc statistics

Displays the FC statistics.

Syntax `show fc statistics {vfabric vfabric-ID | interface fibrechannel}`

Parameters

- *vfabric-ID* — Enter the vfabric ID.
- *fibrechannel* — Enter the Fibre Channel interface name.

Default Not configured

Command Mode EXEC

Usage Information None

Example (vfabric)

```
OS10# show fc statistics vfabric 100
Number of FLOGI                          : 43
Number of FDISC                          : 6
Number of FLOGO                          : 0
Number of FLOGI Accepts                  : 43
Number of FLOGI Rejects                  : 0
Number of FDISC Accepts                  : 6
Number of FDISC Rejects                  : 0
Number of FLOGO Accepts                  : 0
Number of FLOGO Rejects                  : 0
```

Example (interface)

```
OS10# show fc statistics interface fibrechannel1/1/25:1
Number of FLOGI                          : 1
Number of FDISC                          : 0
Number of FLOGO                          : 0
Number of FLOGI Accepts                  : 1
Number of FLOGI Rejects                  : 0
Number of FDISC Accepts                  : 0
Number of FDISC Rejects                  : 0
Number of FLOGO Accepts                  : 0
Number of FLOGO Rejects                  : 0
```

Supported Releases 10.3.1E or later

show fc switch

Displays FC switch parameters.

Syntax `show fc switch`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show fc switch
Switch Mode : FPORT
Switch WWN  : 10:00:14:18:77:20:8d:cf
```

Supported Releases 10.3.1E or later

show running-config vfabric

Displays the running configuration for the vfabric.

Syntax `show running-config vfabric`

Parameters None

Defaults Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show running-configuration vfabric
!
 vfabric 10
 vlan 100
 fcoe fcmap 0xEFC00
 fcoe fcf-priority 140
 fcoe fka-adv-period 13
```

Supported Releases 10.4.0E(R1) or later

show vfabric

Displays vfabric details.

Syntax `show vfabric`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show vfabric
Fabric Name                  SAN_FABRIC
Fabric Type                    FPORT
Fabric Id                         10
VlanId                             1001
FC-MAP                           0EFC00
Config-State                    ACTIVE
Oper-State                       UP
=========================================
 Switch Config Parameters

 =========================================
 Domain ID 4
 =========================================
 Switch Zoning Parameters
 =========================================
 Default Zone Mode: Deny
 Active ZoneSet: zoneset5
 =========================================
 Members
    fibrechannel1/1/25
    port-channel10(Eth 1/1/9)
```

Supported Releases	10.3.1E or later

vfabric

Configures a vfabric.

Syntax	`vfabric` *`fabric-ID`*
Parameters	*`fabric-ID`* — Enter the fabric ID, from 1 to 255.
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	Enable the F_Port or NPG feature before configuring a vfabric. You can configure only one vfabric in F_Port mode. The vfabric becomes active only when you configure the vfabric with a valid VLAN and FC map. Do not use spanned VLAN as vfabric VLAN. The `no` version of this command removes the vfabric. You can remove a vfabric only when it is not applied on any interface.
Example	`OS10(config)# vfabric 100`
Supported Releases	10.3.1E or later

vfabric (interface)

Applies an existing vfabric to an Ethernet or FC interface.

Syntax	`vfabric` *`fabric-ID`*
Parameters	*`fabric-ID`* — Enter the fabric ID, from 1 to 255.
Defaults	Not configured
Command Mode	INTERFACE
Usage Information	The `no` version of this command removes the vfabric from the interface.

Example	`OS10(config)# interface fibrechannel 1/1/1`
	`OS10(conf-if-fc1/1/1)# vfabric 100`
	
	`OS10(config)# interface ethernet 1/1/10`
	`OS10(conf-if-eth1/1/10)# vfabric 200`

Supported Releases	10.3.1E or later

vlan

Associates an existing VLAN ID to the vfabric to carry traffic.

Syntax	`vlan` *`vlan-ID`*
Parameters	*vlan-ID* — Enter an existing VLAN ID.

Defaults	Not configured
Command Mode	Vfabric CONFIGURATION
Usage Information	Create the VLAN ID before associating it to the vfabric. Do not use spanned VLAN as vfabric VLAN. The `no` version of this command removes the VLAN ID from the vfabric.

Example

```
OS10(config)# interface vlan 1023
OS10(conf-if-vl-1023)# exit
OS10(config)# vfabric 100
OS10(conf-vfabric-100)# vlan 1023
```

Supported Releases 10.3.1E or later

FIP-snooping commands

The following commands are supported on FIP-snooping mode:

feature fip-snooping

Enables the FIP snooping feature globally.

Syntax	`feature fip-snooping [with-cvl]`
Parameters	`with-cvl`—To enable CVL.
Defaults	Disabled
Command Mode	CONFIGURATION
Usage Information	You can enable only one of the following at a time: F_Port, NPG, or FSB. You can include the `with-cvl` option to send a Clear Virtual Link (CVL) frame from the FCF to the ENode. This option helps the system to recover automatically if an FCoE session drops. If FIP snooping is already enabled, you can enter the `feature fip-snooping with-cvl` command to enable CVL. You do not have to explicitly disable FIP snooping to enable CVL. However, to disable CVL, you must disable FIP snooping and then re-enable it without the `with-cvl` option. The `no` version of this command disables FIP snooping. When you disable FIP snooping, the system automatically deletes all the FIP snooping VLAN and port mode configurations. If any FIP snooping-related configurations are present in the system, OS10 returns an error message. You can only disable FIP snooping after you remove all the FIP snooping-related configurations from the system.

Example

```
OS10(config)# feature fip-snooping
```

```
OS10(config)# feature fip-snooping with-cvl
```

Supported Releases 10.4.0E(R1) or later

fip-snooping enable

Enables FIP snooping on a specified VLAN.

Syntax `fip-snooping enable`

Parameters	None
Defaults	Disabled
Command Mode	VLAN INTERFACE
Usage Information	Enable FIP snooping on a VLAN only after enabling the FIP snooping feature globally using the `feature fip-snooping` command. OS10 supports FIP snooping on a maximum of 12 VLANs. The `no` version of this command disables FIP snooping on the VLAN.
Example	`OS10(config)# interface vlan 3` `OS10(conf-if-vl-3)# fip-snooping enable`
Supported Releases	10.4.0E(R1) or later

fip-snooping fc-map

Configures the FC map value for a specific VLAN.

Syntax	`fip-snooping fc-map` *`fc-map`*
Parameters	*`fc-map`* — Enter the FC map ID, from `0xefc00` to `0xefcff`.
Defaults	Not configured
Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command disables the FC map configuration.
Example	`OS10(config)# interface vlan 3` `OS10(conf-if-vl-3)# fip-snooping fc-map 0xEFC64`
Supported Releases	10.4.0E(R1) or later

fip-snooping port-mode

Sets FIP snooping port mode for interfaces.

Syntax	`fip-snooping port-mode {enode | enode-transit | fcf | fcf-transit}`
Parameters	`enode | enode-transit | fcf | fcf-transit`—Enter the keyword to set FIP snooping port mode.
Defaults	ENode port mode
Command Mode	INTERFACE
Usage Information	OS10 supports this configuration only on a switch running FSB mode, and on Ethernet and port-channel interfaces. You cannot configure FIP snooping port mode on a port channel member. Use this command to change the port mode. By default, the port mode of an interface is set to ENode. Configure the port mode only after you enable FIP snooping. Before you disable FIP snooping, reset the port mode to its default value, ENode.

You cannot disable FIP snooping when the port mode is set to a non-default value (enode-transit, fcf, or fcf-transit).

If you want to change the port mode from one value to another, you can directly use the `fip-snooping port mode` command. You do not have to explicitly use the `no` form of the command.

The `no` version of this command resets the port mode to ENode.

Example

```
OS10(config)# interface ethernet 1/1/32
OS10(conf-if-eth1/1/32)# fip-snooping port-mode fcf
```

Supported Releases 10.4.0E(R1) or later10.4.3.0 or later—Support for enode-transit and fcf-transit port modes added.

FCoE commands

The following commands are supported on all the three modes: F_Port, NPG, and FSB.

clear fcoe database

Clears the FCoE database for the specified VLAN.

Syntax `clear fcoe database vlan vlan-id {enode enode-mac-address | fcf fcf-mac-address | session fcoe-mac-address}`

Parameters

- *vlan-id* — Enter the VLAN ID.
- *enode-mac-address* — Enter the MAC address of the ENode.
- *fcf-mac-address* — Enter the MAC address of the FCF.
- *fcoe-mac-address* — Enter the MAC address of the FCoE session.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# clear fcoe database vlan 100 enode aa:bb:cc:00:00:00
```

Supported Releases 10.4.0E(R1) or later

clear fcoe statistics

Clears FCoE statistics for specified interface.

Syntax `clear fcoe statistics [interface interface-type]`

Parameters *interface-type* — (Optional) Enter the interface type. The interface may be ethernet, VLAN, or port-channel.

Default Not configured

Command Mode EXEC

Usage Information If you do not specify the `interface interface-type` information, the command clears the statistics for all the interfaces and VLANs.

Example

```
OS10# clear fcoe statistics interface ethernet 1/1/1
OS10# clear fcoe statistics interface port-channel 5
```

Supported Releases 10.4.0E(R1) or later

fcoe-pinned-port

Marks a port as a pinned port in the port-channel. This configuration is supported on FSB, Ethernet LAG in NPG, and F_Port mode. It is not supported on a VLTi LAG.

Syntax `fcoe-pinned-port`

Parameters *node/slot/port[:subport]*—Enter the interface type details.

Defaults Disabled

Command Mode Port-channel INTERFACE

Usage Information You can configure only single port per port-channel. If the port is not configured properly, or if the pinned port goes down, the other ports in the port-channel are not used even if the ports have valid path to server. The `no` version of this command removes the pinned port configuration.

Example

```
OS10(conf-if-eth-1/1/9)# channel-member 10
OS10(conf-if-eth-1/1/9)# fcoe-pinned-port
Warning: Any  existing FCoE session in port-channel will get cleared. Do you
want to continue(yes/no)?yes
```

Supported Releases 10.4.2.0 or later

fcoe max-sessions-per-enodemac

Configures the maximum number of sessions allowed for an ENode.

Syntax `fcoe max-sessions-per-enodemac` *max-session-number*

Parameters *max-session-number* — Enter the maximum number of sessions to be allowed, from 1 to 64.

Defaults 32

Command Mode CONFIGURATION

Usage Information The `no` version of this command resets the number of sessions to the default value.

Example

```
OS10(config)# fcoe max-sessions-per-enodemac 64
```

Supported Releases 10.4.0E(R1) or later

fcoe priority-bits

Configures the priority bits for FCoE application TLVs.

Syntax `fcoe priority-bits` *priority-value*

Parameter *priority-value* — Enter PFC priority value advertised in FCoE application TLV. You can enter one of the following values: `0x01, 0x02, 0x04, 0x08, 0x10, 0x20, 0x40, or 0x80`.

Default	0x08
Command Mode	CONFIGURATION
Usage Information	You can configure only one PFC priority at a time. The `no` version of this command returns the configuration to default value.
Example	`OS10(config)# fcoe priority-bits 0x08`
Supported Releases	10.4.0E(R3) or later

lldp tlv-select dcbxp-appln fcoe

Enables FCoE application TLV for an interface.

Syntax	`lldp tlv-select dcbxp-appln fcoe`
Parameter	None
Default	Enabled
Command Mode	INTERFACE
Usage Information	The default priority value advertised in FCoE application TLV is 3. If the PFC configuration in an interface matches 3, then the FCoE application TLV is advertised as 3. Otherwise, FCoE application TLV is not advertised. When you configure the application priority using `fcoe priority-bits` command, the configured value is advertised in the TLV, which is not dependent on PFC configuration. The `no` version of this command disables the FCoE application TLV.
Example	`OS10(conf-if-eth1/1/1)# lldp tlv-select dcbx-appln fcoe`
Supported Releases	10.4.0E(R3) or later

show fcoe enode

Displays the details of ENodes connected to the switch.

Syntax	`show fcoe enode [`*`enode-mac-address`*`]`
Parameters	*`enode-mac-address`* — (Optional) Enter the MAC address of ENode. This option displays details pertaining to the specified ENode.
Default	Not configured
Command Mode	EXEC
Usage Information	None

Example	`OS10# show fcoe enode`
	`Enode MAC  Enode Interface  VLAN FCFs Sessions`
	`----------  ---------- ---- ---- --------`
	`d4:ae:52:1b:e3:cd  Po 20(Eth 1/1/3)  1001  1  1`

Supported Releases	10.4.0E(R1) or later

show fcoe fcf

Displays details of the FCFs connected to the switch.

Syntax `show fcoe [fcf-mac-address]`

Parameters *fcf-mac-address* — (Optional) Enter the MAC address of the FCF. This option displays details of the specified FCF.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show fcoe fcf
FCF MAC           FCF Interface    VLAN    FC-MAP   FKA_ADV_PERIOD  No. of
Enodes FCF Mode
--------          -------------    ------  -------  --------------
------------ --------
00:0c:84:a8:00:00 Eth 1/1/36        777    0e:fc:00 8000
0             F

00:0d:84:a8:01:02 Eth 1/1/37        778    0e:fc:01 8000
0             FT
```

Supported Releases 10.4.0E(R1) or later

show fcoe pinned-port

Displays the port-channel, the corresponding pinned-port configuration, and the port status if the FCoE sessions are formed.

Syntax `show fcoe pinned-port [port-channel port-channel-id]`

Parameters *port-channel-id*—Enter the port-channel ID to display the corresponding configuration.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show fcoe pinned-port

Interface pinned-port FCoE Status

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

Po 10 Eth 1/1/1 Up

Po 20 Eth 1/1/3 Up

Po 30 Eth 1/1/7 Down
```

Supported Releases 10.4.2.0 or later

show fcoe sessions

Displays the details of the established FCoE sessions.

Syntax `show fcoe sessions [interface vlan vlan-id]`

Parameters *vlan-id* — (Optional) Enter the VLAN ID. This option displays the sessions established on the specified VLAN.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
Enode MAC          Enode Interface  FCF MAC           FCF interface    VLAN
FCoE MAC           FC-ID    PORT WWPN
aa:bb:cc:00:00:00  Po 20(Eth 1/1/3) aa:bb:cd:00:00:00 Po 10(Eth 1/1/1) 100
0e:fc:00:01:00:01  01:00:01 31:00:0e:fc:00:00:00:00 21:00:0e:fc:00:00:00:00
aa:bb:cc:00:00:00  Po 20(Eth 1/1/3) aa:bb:cd:00:00:00 Po 10(Eth 1/1/1) 100
0e:fc:00:01:00:02  01:00:02 31:00:0e:fc:00:00:00:00 21:00:0e:fc:00:00:00:00
```

Supported Releases 10.4.0E(R1) or later

show fcoe statistics

Displays the statistical details of the FCoE control plane.

Syntax `show fcoe statistics [interface interface-type]`

Parameters *interface-type* — (Optional) Enter the type of interface. This option displays statistics of the specified interface.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show fcoe statistics interface port-channel10
Number of Vlan Requests                     :0
Number of Vlan Notifications                :0
Number of Multicast Discovery Solicits      :2
Number of Unicast Discovery Solicits        :0
Number of FLOGI                             :2
Number of FDISC                             :16
Number of FLOGO                             :0
Number of Enode Keep Alive                  :9021
Number of VN Port Keep Alive                :3349
Number of Multicast Discovery Advertisement :4437
Number of Unicast Discovery Advertisement   :2
Number of FLOGI Accepts                     :2
Number of FLOGI Rejects                     :0
Number of FDISC Accepts                     :16
Number of FDISC Rejects                     :0
Number of FLOGO Accepts                     :0
Number of FLOGO Rejects                     :0
Number of CVL                               :0
Number of FCF Discovery Timeouts            :0
Number of VN Port Session Timeouts          :0
Number of Session failures due to Hardware Config :0
```

Supported Releases 10.4.0E(R1) or later

show fcoe system

Displays system information related to the FCoE.

Syntax `show fcoe system`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show fcoe system
Mode: FIP Snooping Bridge
CVL Status: Enabled
FCOE VLAN List (Operational) : 1, 100
FCFs                         : 1
Enodes                       : 2
Sessions                     : 17
```

Supported Releases 10.4.0E(R1) or later

show fcoe vlan

Displays details of FIP-snooping VLANs.

Syntax `show fcoe vlan`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show fcoe vlan
* = Default VLAN
VLAN FC-MAP  FCFs Enodes Sessions
---- ------  ---- ------ --------
*1   -       -    -      -
100 0X0EFC00 1    2      17
```

Supported Releases 10.4.0E(R1) or later

4

Layer 2

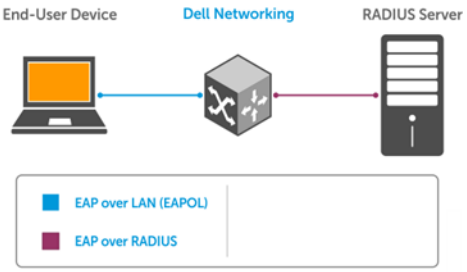
802.1X	Verifies device credentials before sending or receiving packets using the Extensible Authentication Protocol (EAP), see 802.1X Commands.
Link Aggregation Control Protocol (LACP)	Exchanges information between two systems and automatically establishes a link aggregation group (LAG) between the systems, see LACP Commands.
Link Layer Discovery Protocol (LLDP)	Enables a local area network (LAN) device to advertise its configuration and receive configuration information from adjacent LLDP-enabled infrastructure devices, see LLDP Commands.
Media Access Control (MAC)	Configures limits, redundancy, balancing, and failure detection settings for devices on your network using tables, see MAC Commands.
Multiple Spanning-Tree (MST)	Maps MST instances and maps many virtual local area networks (VLANs) to a single spanning-tree instance, reducing the number of required instances, see MST Commands.
Rapid Per-VLAN Spanning-Tree Plus (RPVST+)	Combination of rapid spanning-tree and per-VLAN spanning-tree plus for faster convergence and interoperability, see RPVST+ Commands.
Rapid Spanning-Tree Protocol (RSTP)	Faster convergence and interoperability with devices configured with the Spanning-Tree and Multiple Spanning-Tree Protocols (STPs and MSTPs), see RSTP Commands.
Virtual LANs (VLANs)	Improved security to isolate groups of users into different VLANs and the ability to create a single VLAN across multiple devices, see VLAN Commands.
Port Monitoring (Local/Remote)	Port monitoring of ingress or egress traffic, or both ingress and egress traffic, on specified port(s). Monitoring methods include port-mirroring, remote port monitoring, and encapsulated remote-port monitoring (see Local/Remote Commands).

802.1X

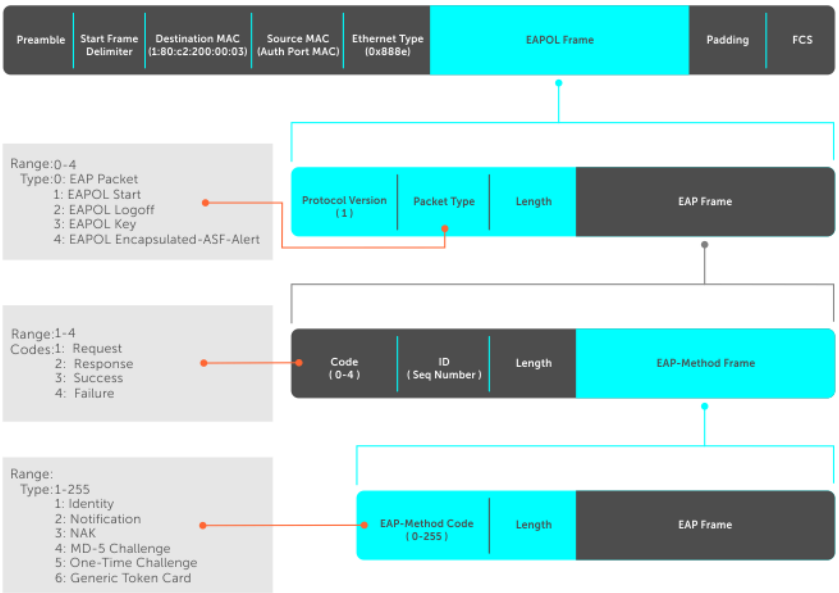
The IEEE 802.1X standard defines a client and server-based access control that prevents unauthorized clients from connecting to a LAN through publicly accessible ports. Authentication is only required in OS10 for inbound traffic. Outbound traffic transmits regardless of the authentication state.

802.1X employs the extensible authentication protocol (EAP) to provide device credentials to an authentication server, typically remote authentication dial-in service (RADIUS), using an intermediary network access device. The network access device mediates all communication between the end-user device and the authentication server so the network remains secure.

The network access device uses EAP-over-Ethernet, also known as EAPOL — EAP over LAN, to communicate with the end user device and EAP-over-RADIUS to communicate with the server.



NOTE: OS10 supports only RADIUS as the back-end authentication server.



The authentication process involves three devices:

- **Supplicant** — The device attempting to access the network performs the role of supplicant. Regular traffic from this device does not reach the network until the port associated to the device is authorized. Before that, the supplicant can only exchange 802.1x messages (EAPOL frames) with the authenticator.
- **Authenticator** — The authenticator is the gate keeper of the network, translating and forwarding requests and responses between the authentication server and the supplicant. The authenticator also changes the status of the port based on the results of the authentication process. The authenticator executes on the Dell EMC device.
- **Authentication-server** — The authentication-server selects the authentication method, verifies the information the supplicant provides, and grants network access privileges.

Port authentication

The process begins when the authenticator senses a link status change from down to up:

1 The authenticator requests that the supplicant identify itself using an EAP *Request Identity* frame.

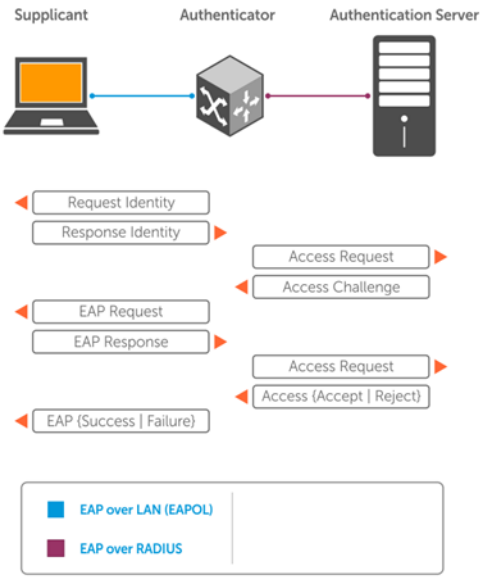
2 The supplicant responds with its identity in an EAP *Response Identity* frame.

3 The authenticator decapsulates the EAP response from the EAPOL frame, encapsulates it in a RADIUS *Access Request* frame, and forwards the frame to the authentication server.

4 The authentication server replies with an *Access Challenge* frame who requests that the supplicant verifies its identity using an EAP-Method. The authenticator translates and forwards the challenge to the supplicant.

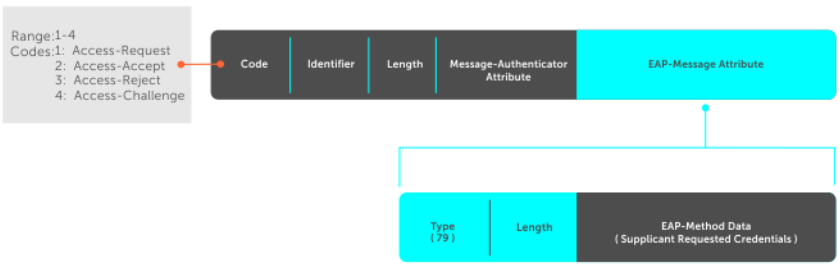
5 The supplicant negotiates the authentication method and provides the *EAP Request* information in an *EAP Response*. Another *Access Request* frame translates and forwards the response to the authentication server.

6 If the identity information the supplicant provides is valid, the authentication server sends an *Access Accept* frame that specify the network privileges. The authenticator changes the port state to authorize and forwards an *EAP Success* frame. If the identity information is invalid, the server sends an *Access Reject* frame. If the port state remains unauthorized, the authenticator forwards an *EAP Failure* frame.



EAP over RADIUS

802.1X uses RADIUS to transfer EAP packets between the authenticator and the authentication server. EAP messages are encapsulated in RADIUS packets as an attribute of type, length, value (TLV) format — the *type* value for EAP messages is 79.

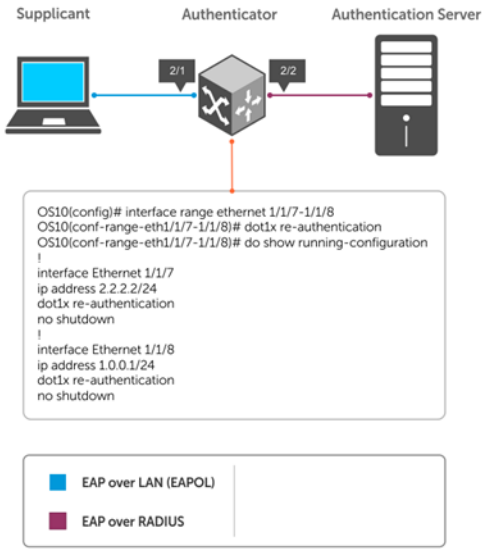


Configure 802.1X

You can configure and enable 802.1X on a port in a single process. OS10 supports 802.1X with EAP-MD5, EAP-OTP, EAP-TLS, EAP-TTLS, PEAPv0, PEAPv1, and MS-CHAPv2 with PEAP. All platforms support RADIUS as the authentication server.

If the primary RADIUS server becomes unresponsive, the authenticator begins using a secondary RADIUS server if configured.

NOTE: 802.1X is not supported on port-channels or port-channel members.



Enable 802.1X

1 Enable 802.1X globally in CONFIGURATION mode.

```
dot1x system-auth-control
```

2 Enter an interface or a range of interfaces in INTERFACE mode.

```
interface range
```

3 Enable 802.1X on the supplicant interface only in INTERFACE mode.

```
dot1x port-control auto
```

Configure and verify 802.1X configuration

```
OS10(config)# dot1x system-auth-control
OS10(config)# interface range 1/1/7-1/1/8
OS10(conf-range-eth1/1/7-1/1/8)# dot1x port-control auto
OS10(conf-range-eth1/1/7-1/1/8)# dot1x re-authentication
OS10(conf-range-eth1/1/7-1/1/8)# do show dot1x interface ethernet 1/1/7

802.1x information on ethernet1/1/7
-------------------------------------
Dot1x Status:              Enable
Port Control:              AUTO
Port Auth Status:          UNAUTHORIZED
Re-Authentication:         Enable
Tx Period:                 60 seconds
Quiet Period:              60 seconds
Supplicant Timeout:        30 seconds
Server Timeout:            30 seconds
Re-Auth Interval:          3600 seconds
Max-EAP-Req:               2
Host Mode:                 MULTI_HOST
Auth PAE State:            Initialize
Backend State:             Idle
```

Identity retransmissions

If the authenticator sends a *Request Identity* frame but the supplicant does not respond, the authenticator waits 30 seconds and then retransmits the frame. There are several reasons why the supplicant might fail to respond — the supplicant maybe booting when the request arrived, there may be a physical layer problem, and so on.

1 Configure the amount of time that the authenticator waits before retransmitting an EAP *Request Identity* frame in INTERFACE mode, from 1 to 65535 – 1 year, default 60.

```
dot1x timeout tx-period seconds
```

2 Configure a maximum number of times the authenticator retransmits a *Request Identity* frame in INTERFACE mode from 1 to 10, default 2.

```
dot1x max-req retry-count
```

Configure and verify retransmission time

```
OS10(config)# dot1x system-auth-control
OS10(config)# interface range 1/1/7-1/1/8
OS10(conf-range-eth1/1/7-1/1/8)# dot1x timeout tx-period 120
OS10(conf-range-eth1/1/7-1/1/8)# dot1x max-req 5
OS10(conf-range-eth1/1/7-1/1/8)# do show dot1x interface ethernet 1/1/7

802.1x information on ethernet1/1/7
-----------------------------------
Dot1x Status:              Enable
Port Control:              AUTO
Port Auth Status:          UNAUTHORIZED
Re-Authentication:         Enable
Tx Period:                 120 seconds
Quiet Period:              60 seconds
Supplicant Timeout:        30 seconds
Server Timeout:            30 seconds
Re-Auth Interval:          3600 seconds
Max-EAP-Req:               5
Host Mode:                 MULTI_HOST
Auth PAE State:            Initialize
Backend State:             Idle
```

View interface running configuration

```
OS10(conf-range-eth1/1/7-1/1/8)# do show running-configuration interface
...
!
interface ethernet1/1/7
 no shutdown
 dot1x max-req 5
 dot1x port-control auto
 dot1x re-authentication
 dot1x timeout quiet-period 120
 dot1x timeout tx-period 120
!
interface ethernet1/1/8
 no shutdown
 dot1x max-req 5
 dot1x port-control auto
 dot1x re-authentication
 dot1x timeout quiet-period 120
 dot1x timeout tx-period 120
...
```

Failure quiet period

If the supplicant fails the authentication process, the authenticator sends another Request Identity frame after 30 seconds by default. The quiet period is a transmit interval time after a failed authentication.

The Request Identity Retransmit interval is for an unresponsive supplicant. You can configure the interval for a maximum of 10 times for an unresponsive supplicant.

1 Configure the amount of time that the authenticator waits to retransmit a *Request Identity* frame after a failed authentication in INTERFACE mode from 1 to 65535, default 60 seconds.

```
dot1x timeout quiet-period seconds
```

Configure and verify port authentication

```
OS10(conf-range-eth1/1/7-1/1/8)# dot1x timeout quiet-period 120
OS10(conf-range-eth1/1/7-1/1/8)# do show dot1x interface ethernet 1/1/7
802.1x information on ethernet1/1/7
-----------------------------------
Dot1x Status:              Enable
Port Control:              AUTO
Port Auth Status:          UNAUTHORIZED
Re-Authentication:         Enable
Tx Period:                 120 seconds
Quiet Period:              120 seconds
Supplicant Timeout:        30 seconds
Server Timeout:            30 seconds
Re-Auth Interval:          3600 seconds
Max-EAP-Req:               5
Host Mode:                 MULTI_HOST
Auth PAE State:            Initialize
Backend State:             Idle
```

View interface running configuration

```
OS10(conf-range-eth1/1/7-1/1/8)# do show running-configuration interface
...
!
interface ethernet1/1/7
 no shutdown
 dot1x max-req 5
 dot1x port-control auto
 dot1x re-authentication
 dot1x timeout quiet-period 120
 dot1x timeout tx-period 120
!
interface ethernet1/1/8
 no shutdown
 dot1x max-req 5
 dot1x port-control auto
 dot1x re-authentication
 dot1x timeout quiet-period 120
 dot1x timeout tx-period 120
...
```

Port control mode

802.1X requires a port to be in one of three states — force-authorized, force-unauthorized, or auto.

force-authorized (default)	This is an *authorized state*. A device connected to this port does not use the authentication process but can communicate on the network. Placing the port in this state is the same as disabling 802.1X on the port. `force-authorized` is the default mode.
force-unauthorized	This is an *unauthorized state*. A device connected to a port does not use the authentication process but is *not* allowed to communicate on the network. Placing the port in this state is the same as shutting down the port. Any attempt by the supplicant to initiate authentication is ignored.
auto	This is an *unauthorized state* by default. A device connected to this port is subject to the authentication process. If the process is successful, the port is authorized and the connected device communicates on the network.

- Place a port in the auto, force-authorized (default), or force-unauthorized state in INTERFACE mode.
  ```
  dot1x port-control {auto | force-authorized | force-unauthorized}
  ```

Configure and verify force-authorized state

```
OS10(conf-range-eth1/1/7-1/1/8)# dot1x port-control force-authorized
OS10(conf-range-eth1/1/7-1/1/8)# do show dot1x interface ethernet 1/1/7

802.1x information on ethernet1/1/7
-----------------------------------
Dot1x Status:              Enable
Port Control:              AUTHORIZED
Port Auth Status:          UNAUTHORIZED
Re-Authentication:         Enable
Tx Period:                 120 seconds
Quiet Period:              120 seconds
Supplicant Timeout:        30 seconds
Server Timeout:            30 seconds
Re-Auth Interval:          3600 seconds
Max-EAP-Req:               5
Host Mode:                 MULTI_HOST
Auth PAE State:            Initialize
Backend State:             Initialize
```

View interface running configuration

```
OS10(conf-range-eth1/1/7-1/1/8)# do show running-configuration interface
...
!
interface ethernet1/1/7
 no shutdown
 dot1x max-req 5
 dot1x re-authentication
 dot1x timeout quiet-period 120
 dot1x timeout tx-period 120
!
interface ethernet1/1/8
 no shutdown
 dot1x max-req 5
 dot1x re-authentication
 dot1x timeout quiet-period 120
 dot1x timeout tx-period 120
...
```

Reauthenticate port

Configures the time period for reauthentication. After the supplicant is authenticated and the port is authorized, configure the authenticator to reauthenticate the supplicant. If you enable reauthentication, the supplicant reauthenticates every 3600 seconds.

- Re-authenticate the supplicant in INTERFACE mode, from 1 to 65535, default 3600.
  ```
  dot1x timeout re-authperiod seconds
  ```

Configure and verify reauthentication time period

```
OS10(config)# interface range ethernet 1/1/7-1/1/8
OS10(conf-range-eth1/1/7-1/1/8)# dot1x re-authentication
OS10(conf-range-eth1/1/7-1/1/8)# dot1x timeout re-authperiod 3600
OS10(conf-range-eth1/1/7-1/1/8)# show dot1x interface ethernet 1/1/7

802.1x information on ethernet1/1/7
-----------------------------------
Dot1x Status:              Enable
Port Control:              AUTHORIZED
Port Auth Status:          UNAUTHORIZED
Re-Authentication:         Enable
Tx Period:                 120 seconds
Quiet Period:              120 seconds
Supplicant Timeout:        30 seconds
Server Timeout:            30 seconds
Re-Auth Interval:          3600 seconds
Max-EAP-Req:               5
Host Mode:                 MULTI_HOST
Auth PAE State:            Initialize
Backend State:             Initialize
```

View interface running configuration

```
OS10(conf-range-eth1/1/7-1/1/8)# do show running-configuration interface
...
!
interface ethernet1/1/7
 no shutdown
 dot1x max-req 5
 dot1x re-authentication
 dot1x timeout quiet-period 120
 dot1x timeout re-authperiod 3600
 dot1x timeout tx-period 120
!
interface ethernet1/1/8
 no shutdown
 dot1x max-req 5
 dot1x re-authentication
 dot1x timeout quiet-period 120
 dot1x timeout re-authperiod 3600
 dot1x timeout tx-period 120
...
```

Configure timeouts

If the supplicant or authentication server is unresponsive, the authenticator terminates the authentication process after 30 seconds by default. Configure the amount of time the authenticator waits for a response before termination.

- Terminate the authentication process due to an unresponsive supplicant in INTERFACE mode, from 1 to 65535, default 30.
  `dot1x timeout supp-timeout seconds`
- Terminate the authentication process due to an unresponsive authentication server in INTERFACE mode, from 1 to 65535, default 30.
  `dot1x timeout server-timeout seconds`

Configure and verify server timeouts

```
OS10(conf-range-eth1/1/7-1/1/8)# dot1x timeout supp-timeout 45
OS10(conf-range-eth1/1/7-1/1/8)# dot1x timeout server-timeout 60
OS10(conf-range-eth1/1/7-1/1/8)# do show dot1x interface ethernet 1/1/7

802.1x information on ethernet1/1/7
-----------------------------------
Dot1x Status:              Enable
Port Control:              AUTHORIZED
```

```
Port Auth Status:          UNAUTHORIZED
Re-Authentication:         Enable
Tx Period:                 120 seconds
Quiet Period:              120 seconds
Supplicant Timeout:        45 seconds
Server Timeout:            60 seconds
Re-Auth Interval:          3600 seconds
Max-EAP-Req:               5
Host Mode:                 MULTI_HOST
Auth PAE State:            Initialize
Backend State:             Initialize
```

View interface running configuration

```
OS10(conf-range-eth1/1/7-1/1/8)# do show running-configuration interface
...
!
interface ethernet1/1/7
 no shutdown
 dot1x max-req 5
 dot1x re-authentication
 dot1x timeout quiet-period 120
 dot1x timeout re-authperiod 3600
 dot1x timeout server-timeout 60
 dot1x timeout supp-timeout 45
 dot1x timeout tx-period 120
!
interface ethernet1/1/8
 no shutdown
 dot1x max-req 5
 dot1x re-authentication
 dot1x timeout quiet-period 120
 dot1x timeout re-authperiod 3600
 dot1x timeout server-timeout 60
 dot1x timeout supp-timeout 45
 dot1x timeout tx-period 120
...
```

802.1X commands

dot1x host-mode

Allows 802.1X authentication for either a single supplicant or multiple supplicants on an interface.

Syntax `dot1x host-mode {multi-host}`

Parameters

- `multi-host` — Allows attachment of multiple hosts to a single 802.1X-enabled port. You can only authorize one of the attached clients for all clients to grant network access. If the port becomes unauthorized (re-authentication fails or receives an EAPOL-logoff message), the device denies network access to all of the attached clients.

Default Multi-host

Command Mode INTERFACE

Usage Information The `no` version of this command resets the value to the default.

Example

```
OS10(conf-range-eth1/1/7-1/1/8)# dot1x host-mode multi-host
```

Supported Releases 10.2.0E or later

dot1x max-req

Changes the maximum number of requests that the device sends to a supplicant before restarting 802.1X authentication.

Syntax `dot1x max-req retry-count`

Parameters `max-req retry-count` — Enter the retry count for the request sent to the supplicant before restarting 802.1X reauthentication, from 1 to 10.

Default 2

Command Mode INTERFACE

Usage Information The `no` version of this command resets the value to the default.

Example

```
OS10(conf-range-eth1/1/7-1/1/8)# dot1x max-req 4
```

Supported Releases 10.2.0E or later

dot1x port-control

Controls the 802.1X authentication performed on the interface.

Syntax `dot1x port-control {force-authorized | force-unauthorized | auto}`

Parameters

- `force-authorized` — Disables 802.1X authentication on the interface and allows all traffic on the interface without authentication.
- `force-unauthorized` — Keeps the port in the unauthorized state, ignoring all attempts by the client to authenticate.
- `auto` — Enables 802.1X authentication on the interface.

Default Force-authorized

Command Mode INTERFACE

Usage Information The `no` version of this command resets the value to the default.

Example

```
OS10(config)# interface range ethernet 1/1/7-1/1/8
OS10(conf-range-eth1/1/7-1/1/8)# dot1x port-control auto
```

Supported Releases 10.2.0E or later

dot1x re-authentication

Enables periodic re-authentication of 802.1X supplicants.

Syntax `dot1x re-authentication`

Parameters None

Default Disabled

Command Mode INTERFACE

Usage Information The `no` version of this command disables the periodic re-authentication of 8021.X supplicants.

Example	`OS10(conf-range-eth1/1/7-1/1/8)# dot1x re-authentication`
Supported Releases	10.2.0E or later

dot1x timeout quiet-period

Sets the number of seconds that the device remains in the quiet state following a failed authentication exchange with a supplicant.

Syntax	`dot1x timeout quiet-period` *seconds*
Parameters	`quiet period` *seconds* — Enter the number of seconds for the 802.1X quiet period timeout, from 1 to 65535.
Default	60 seconds
Command Mode	INTERFACE
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(conf-range-eth1/1/7-1/1/8)# dot1x timeout quiet-period 120`
Supported Releases	10.2.0E or later

dot1x timeout re-authperiod

Sets the number of seconds between re-authentication attempts.

Syntax	`dot1x timeout re-authperiod` *seconds*
Parameters	`re-authperiod` *seconds* — Enter the number of seconds for the 802.1X re-authentication timeout, from 1 to 65535.
Default	3600 seconds
Command Mode	INTERFACE
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(conf-range-eth1/1/7-1/1/8)# dot1x timeout re-authperiod 7200`
Supported Releases	10.2.0E or later

dot1x timeout server-timeout

Sets the number of seconds that the device waits before retransmitting a packet to the authentication server.

Syntax	`dot1x timeout server-timeout` *seconds*
Parameters	`server-timeout` *seconds* — Enter the number of seconds for the 802.1X server timeout, from 1 to 65535.
Default	30 seconds
Command Mode	INTERFACE
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(conf-range-eth1/1/7-1/1/8)# dot1x server-timeout 60`

Supported Releases	10.2.0E or later

dot1x timeout supp-timeout

Sets the number of seconds that the device waits for the supplicant to respond to an EAP request frame before the device retransmits the frame.

Syntax	`dot1x timeout supp-timeout` *seconds*
Parameters	`supp-timeout` *seconds* — Enter the number of seconds for the 802.1X supplicant timeout, from 1 to 65535.
Default	30 seconds
Command Mode	INTERFACE
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(conf-range-eth1/1/7-1/1/8)# dot1x timeout supp-timeout 45`
Supported Releases	10.2.0E or later

dot1x timeout tx-period

Sets the number of seconds that the device waits for a response to an EAP-request/identity frame from the supplicant before retransmitting the request.

Syntax	`dot1x timeout tx-period` *seconds*
Parameters	`tx-period` *seconds* — Enter the number of seconds for the 802.1X transmission timeout, from 1 to 65535.
Default	60 seconds
Command Mode	INTERFACE
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(conf-range-eth1/1/7-1/1/8)# dot1x timeout tx-period 120`
Supported Releases	10.2.0E or later

show dot1x

Displays global 802.1X configuration information.

Syntax	`show dot1x`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	None

Example	`OS10# show dot1x`
	`PAE Capability: Authenticator only`
	`Protocol Version: 2`
	`System Auth Control: Enable`
	`Auth Server: Radius`

Supported Releases 10.2.0E or later

show dot1x interface

Displays 802.1X configuration information.

Syntax `show dot1x interface ethernet node/slot/port[:subport]`

Parameters `ethernet node/slot/port[:subport]` — Enter the Ethernet interface information.

Command Mode EXEC

Usage Information Use this command to view the dot1x interface configuration for a specific interface.

Example

```
OS10# show dot1x interface
802.1x information on ethernet1/1/1
-------------------------------------
Dot1x Status:              Enable
802.1x information on ethernet1/1/2
-------------------------------------
Dot1x Status:              Enable
802.1x information on ethernet1/1/3
-------------------------------------
Dot1x Status:              Enable
802.1x information on ethernet1/1/4
-------------------------------------
Dot1x Status:              Enable
802.1x information on ethernet1/1/5
-------------------------------------
Dot1x Status:              Enable
802.1x information on ethernet1/1/6
-------------------------------------
Dot1x Status:              Enable
802.1x information on ethernet1/1/7
-------------------------------------
Dot1x Status:              Enable
Port Control:              AUTO
Port Auth Status:          UNAUTHORIZED
--more--
```

Example (when dot1x is not enabled globally)

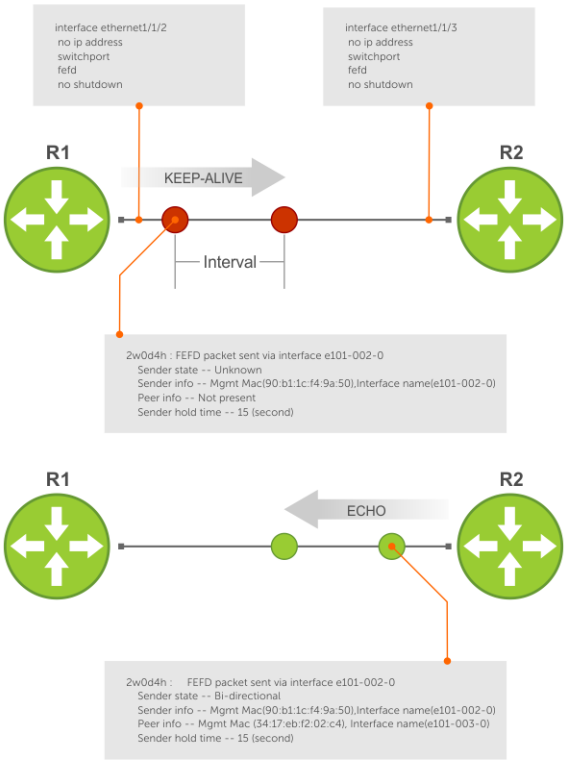
```
OS10# show dot1x interface
802.1x not enabled in the system
OS10#
```

Supported Releases 10.2.0E or later

Far-end failure detection

Far-End Failure Detection (FEFD) is a protocol that senses remote data link errors in a network.

FEFD uses a link layer echo protocol to detect and signal far-end failures over Ethernet and optical links. When you enable FEFD, switches periodically exchange FEFD echo Ethernet frames to identify link failures. If the local switch does not receive an echo frame from its peer for three times the FEFD message interval, the local switch assumes that the peer link is down.



FEFD helps detect far-end failure when the following problems occur:

- Only one side receives packets although the physical Layer 1 (L1) link is up on both sides.
- Transceivers are not connected to the correct ports.

FEFD states

FEFD comprises the following four states:

- `Idle`—FEFD is disabled.
- `Unknown`—Shown when FEFD is enabled and changes to `bi-directional` after successful handshake with the peer. Also shown if the peer goes down in `normal` mode.
- `bi-directional`—Interface is up, connected, and receiving echoes of its neighbor.
- `err-disabled`—Only found when FEFD mode is `aggressive` and when the interface has not received three echoes of its neighbor. To reset an interface in this state, use the `fefd reset` command.

FEFD modes

FEFD operates in two modes—normal mode and aggressive mode.

- Normal mode—When you enable Normal mode on an interface and a far-end failure is detected, no intervention is required to reset the interface to bring it back to an FEFD operational state.
- Aggressive mode—When you enable Aggressive mode on an interface in the same state, you must manually reset the interface.

The following events explain how FEFD state transition occurs:

- When you enable FEFD on an interface a link transitions from idle state to unknown state.
- In the unknown state, the interface starts transmitting link state information at a regular interval. The interface state changes to bi-directional when a handshake is complete with the peer.
- If an interface does not receive three echo messages from its peer when an interface is in bi-directional state, the interface state changes to unknown in normal mode. In aggressive mode, the interface changes to err-disabled state.

If the interface state changes to err-disabled, use the `fefd reset [interface]` global command to reset these interfaces. The unknown or err-disabled state brings the line protocol down so that the upper protocols can detect that the peer link is down.

Table 4. FEFD state changes

Local event (User intervention)	Mode (In)	Local state (Show display) (Result)	Local admin State (Result)	Local line protocol Status (Result)	Remote state (Show display) (Result)	Remote admin state	Remote line protocol status
Shutdown(user configuration)	Normal	Admin Shutdown	Down	Down	Line protocol is down	Up	Down
Shutdown(user configuration)	Aggressive	Admin Shutdown	Down	Down	Line protocol is down	Up	Down
FEFD+ FEFD disable(user configuration)	Normal	Locally disabled	Up	Up	Unknown	Up	Down
FEFD + FEFD disable(user configuration)	Aggressive	Locally disabled	Up	Down	Err-disabled	Up	Down
Link Failure (Remove cable or transceiver)	Normal	Unknown	Up	Down	Unknown	Up	Down
Link Failure(Remove cable or transceiver)	Aggressive	Unknown	Up	Down	Unknown	Up	Down
FEFD enable(user configuration)	Normal	Bi-directional	Up	Up	Bi-directional	Up	Up
FEFD enable(user configuration)	Aggressive	Bi-directional	Up	Up	Bi-directional	Up	Up

Restrictions

- You can enable FEFD globally or on an interface. The FEFD interface configuration overrides the global FEFD configuration.
- OS10 supports FEFD only on physical Ethernet interfaces. FEFD is not supported on logical interfaces.

Enable FEFD globally

To configure FEFD globally:

1 Do one of the following:

- Configure FEFD normal mode globally using the `fefd-global` command in CONFIGURATION mode.
  ```
  OS10(Config)# fefd-global
  ```
- Configure FEFD normal mode globally using the `fefd-global mode normal` command in CONFIGURATION mode.
  ```
  OS10(Config)# fefd-global mode normal
  ```
- Configure FEFD aggressive mode globally using the `fefd-global mode aggressive` command in CONFIGURATION mode.
  ```
  OS10(Config)# fefd-global mode aggressive
  ```

2 (Optional) Configure the FEFD interval using the `fefd-global interval` command in CONFIGURATION mode and enter the interval in seconds. The range is from 3 to 255 seconds.

```
OS10(Config)# fefd-global interval 20
```

3 (Optional) Disable FEFD on a specific interface if required using the `fefd disable` command in INTERFACE mode.

```
OS10(Conf-if-eth1/1/1)# no fefd interval 20
```

Enable FEFD on interface

To configure FEFD on an interface:

1 Do one of the following:

- Configure FEFD normal mode on an interface using the `fefd` command in INTERFACE mode.

  ```
  OS10(Conf-if-eth1/1/1)# fefd
  ```

- Configure FEFD normal mode on an interface using the `fefd mode normal` command in INTERFACE mode.

  ```
  OS10(Conf-if-eth1/1/1)# fefd mode normal
  ```

- Configure FEFD aggressive mode on an interface using the `fefd mode aggressive` command in INTERFACE mode.

  ```
  OS10(Conf-if-eth1/1/1)# fefd mode aggressive
  ```

2 (Optional) Configure the FEFD interval using the `fefd interval` command in INTERFACE mode and enter the interval in seconds. The range is from 3 to 255 seconds.

```
OS10(Conf-if-eth1/1/1)# fefd interval 20
```

Reset FEFD err-disabled interface

When the system detects a far-end failure in FEFD aggressive mode, the interface moves to err-disabled state. To bring back the interface to FEFD operational state:

- Enter the `fefd reset` command in EXEC mode.

  ```
  OS10# fefd reset ethernet 1/1/1
  ```

Display FEFD information

To view FEFD information:

- To view FEFD information globally, use the `show fefd` command in EXEC mode.
- To view FEFD information for an interface, use the `show fefd interface` command in EXEC mode.

The following is a sample output of FEFD global information:

```
OS10# show fefd
FEFD is globally 'ON', interval is 15 seconds, mode is Normal.
INTERFACE      MODE          INTERVAL      STATE
============================================================
eth1/1/1       NA            NA            Idle (Not running)
eth1/1/2       NA            NA            Idle (Not running)
eth1/1/3       NA            NA            Idle (Not running)
eth1/1/4       NA            NA            Idle (Not running)
eth1/1/5       NA            NA            Idle (Not running)
eth1/1/6       NA            NA            Idle (Not running)
eth1/1/7       NA            NA            Idle (Not running)
```

The following is a sample output of FEFD information for an interface:

```
rt-maa-s4248FBL-3# show fefd ethernet 1/1/1
FEFD is globally 'ON', interval is 15 seconds, mode is Normal.
```

```
INTERFACE       MODE            INTERVAL        STATE
============================================================
eth1/1/1        NA              NA              Idle (Not running)
```

FEFD Commands

debug fefd

Enables debugging of FEFD.

Syntax `debug fefd {all | events | packets} [interface]`

Parameters

- `all`—Enter the keyword to view all FEFD debug information.
- `events`—Enter the keyword to view debug information about FEFD state changes.
- `packets`—Enter the keyword to view debug information about FEFD packets that are sent and received.
- (Optional) *`interface`*—Enter interface information.

Default Not configured

Command Mode EXEC Privilege

Example

```
OS10# debug fefd
```

Supported Releases 10.4.3.0 or later

fefd

Configures FEFD on an interface.

Syntax `fefd [mode {normal | aggressive} | interval seconds | disable]`

Parameters

- (Optional) `mode`—Enter the keyword and enter either `normal` to enable the normal mode or `aggressive` to enable the aggressive mode.
- `interval`—Enter the keyword and enter the FEFD interval in seconds to configure the interval between FEFD control packets on an interface. The range is from 3 to 255. The default value is 15 seconds.
- `disable`—Enter the keyword to disable FEFD on a specific interface when you configure FEFD globally.

Default Not configured

Command Mode INTERFACE

Usage Information The `fefd` command without any arguments enables the normal mode with the default FEFD interval of 15 seconds.

If you use the `no fefd` command, the system does not disable FEFD if the `fefd mode` command is already present in the configuration. Similarly, if you use the `no fefd mode` command, the system does not disable FEFD if the `fefd` command is already present in the configuration.

To disable FEFD on an interface when FEFD globally enabled, use the `fefd disable` command on the interface.

To unconfigure FEFD on an interface, use either the `no fefd` command or the `no fefd mode` command. To return to the default FEFD interval, use the `no fefd interval` command.

Example

```
OS10(conf-if-eth1/1/9)# fefd
OS10(conf-if-eth1/1/9)# fefd mode aggressive
OS10(conf-if-eth1/1/9)# fefd mode interval 10
```

Supported Releases 10.4.3.0 or later

fefd-global

Configures FEFD globally.

Syntax `fefd-global [mode {normal | aggressive} | interval seconds]`

Parameters

- (Optional) `mode`—Enter the keyword and enter either `normal` to enable the normal mode or `aggressive` to enable the aggressive mode.
- (Optional) `interval`—Enter the keyword and enter the FEFD interval in seconds to configure the interval between FEFD control packets globally. The range is from 3 to 255. The default value is 15 seconds.

Default Not configured

Command Mode CONFIGURATION

Usage Information The `fefd-global` command without arguments enables normal mode with the default FEFD interval of 15 sseconds.

If you use the `no fefd-global` command, the system does not disable FEFD if the `fefd-global mode` command is already present in the configuration. Similarly, if you use the `no fefd-global mode` command, the system does not disable FEFD if the `fefd-global` command is already present in the configuration.

To unconfigure FEFD globally, use either the `no fefd-global` command or the `no fefd-global mode` command. To return to the default FEFD interval, use the `no fefd-global interval` command.

Example

```
OS10(config)# fefd-global
OS10(config)# fefd-global mode aggressive
OS10(config)# fefd mode interval 10
```

Supported Releases 10.4.3.0 or later

fefd reset

Resets interfaces that are in error-disabled state because FEFD is set to aggressive mode.

Syntax `fefd reset [interface]`

Parameters

- (Optional) *interface*—Enter the interface name to reset the error-disabled state of the interface because FEFD is set to aggressive mode.

Default Not configured

Command Mode EXEC

Usage Information If you do not enter the interface name, this command resets the error-disabled state of all interfaces because FEFD is set to aggressive mode.

Example

```
OS10# fefd reset
```

```
OS10# fefd reset ethernet 1/1/2
```

Supported Releases 10.4.3.0 or later

show fefd

Displays FEFD information globally or for a specific interface.

Syntax `show fefd [interface]`

Parameters

- (Optional) *interface*—Enter the interface information.

Default Not configured

Command Mode EXEC and EXEC Privilege

Usage Information The following table describes the fields in the `show fefd` command output:

Field	Description
Interface	Displays the interface name
Mode	Displays the mode—aggressive, normal, or NA when the interface contains `fefd reset` in its configuration.
Interval	Displays the interval between FEFD packets.
State	Displays the state of the interface and can be one of the following: • Bi-directional—Interface is up, connected, and receiving echoes of its neighbor. • Err-disabled—Only found when FEFD mode is aggressive and when the interface has not received three echoes of its neighbor. To reset an interface in this state, use the `fefd reset` command. • `Unknown`—Shown when FEFD is enabled and changes to `bi-directional` after successful handshake with the peer. Also shown if the peer goes down in `normal` mode. • Locally disabled—Interface contains the `fefd reset` command in its configuration. • Admin Shutdown—Interface is disabled using the shutdown command. • Line protocol is down—The state on the remote device when an interface of the local device is disabled with the `shutdown` command.

Example

```
OS10# show fefd
FEFD is globally 'ON', interval is 22 seconds,mode is NORMAL.
INTERFACE MODE       INTERVAL STATE
============================================================
eth1/1/1  Normal     22       Unknown
eth1/1/2  Normal     22       Unknown
eth1/1/3  Normal     22       Unknown
eth1/1/4  Normal     22       Unknown
eth1/1/5  Normal     22       Unknown
eth1/1/6  Normal     22       Unknown
eth1/1/7  Normal     22       Unknown
```

```
eth1/1/8  Normal     22        Unknown
eth1/1/9  Aggressive 22        Err-disabled
eth1/1/10 Normal     22        Unknown
```

Supported Releases 10.4.3.0 or later

Link Aggregation Control Protocol

Group Ethernet interfaces to form a single link layer interface called a LAG or port-channel. Aggregating multiple links between physical interfaces creates a single logical LAG, which balances traffic across the member links within an aggregated Ethernet bundle and increases the uplink bandwidth. If one member link fails, the LAG continues to carry traffic over the remaining links. For information about LAG load balancing and hashing, see Load balancing.

You can use LACP to create dynamic LAGs exchanging information between two systems (also called Partner Systems) and automatically establishing the LAG between the systems. LACP permits the exchange of messages on a link to:

- Reach an agreement on the identity of the LAG to which the link belongs.
- Move the link to that LAG.
- Enable the transmission and reception functions.

LACP functions by constantly exchanging custom MAC PDUs across LAN Ethernet links. The protocol only exchanges packets between ports you configure as LACP-capable.

Modes

A LAG includes three configuration modes — on, active, and passive.

On	Sets the Channeling mode to Static. The interface acts as a member of the static LAG.
Active	Sets the interface in the Active Negotiating state. LACP runs on any link configured in this mode. A port in Active mode automatically initiates negotiations with other ports by using LACP packets. A port in Active mode can set up a port-channel (LAG) with another port in Active mode or Passive mode.
Passive	Sets the interface in an Inactive Negotiating state, but LACP runs on the link. A port in Passive mode also responds to negotiation requests (from ports in Active mode). Ports in Passive mode respond to LACP packets. A port in Passive mode cannot set up a LAG with another port in Passive mode.

- There is no dual-membership in static and dynamic LAGs:
  - If a physical interface is a part of a static LAG, the `channel-group id mode active` command is rejected on that interface.
  - If a physical interface is a part of a dynamic LAG, the `channel-group id` command is rejected on that interface.
- You cannot add static and dynamic members to the same LAG.
- There is a difference between the `shutdown` and `no interface port-channel` commands:
  - The `shutdown` command on LAG *xyz* disables the LAG and retains the user commands.
  - The `no interface port-channel channel-number` command deletes the specified LAG, including a dynamically created LAG. The interfaces restore and are ready for configuration.
- A maximum of 128 port-channels with up to 16 members per channel are allowed.

Configuration

LACP is enabled globally by default. You can configure aggregated ports with compatible active and passive LACP modes to automatically link them.

1 Configure the system priority in CONFIGURATION mode (1 to 65535; the higher the number, the lower the priority; default 32768).

```
lacp system-priority priority-value
```

2 Configure the LACP port priority in INTERFACE mode (1 to 65535; the higher the number, the lower the priority; default 32768).

```
lacp port-priority priority-value
```

3 Configure the LACP rate in INTERFACE mode (default normal).

```
lacp rate [fast | normal]
```

Configure LACP

```
OS10(config)# lacp system-priority 65535
OS10(config)# interface range ethernet 1/1/7-1/1/8
OS10(conf-range-eth1/1/7-1/1/8)# lacp port-priority 4096
OS10(conf-range-eth1/1/7-1/1/8)# lacp rate fast
```

Verify LACP configuration

```
OS10(conf-range-eth1/1/7-1/1/8)# do show running-configuration
...
!
interface ethernet1/1/7
 lacp port-priority 4096
 lacp rate fast
 no shutdown
!
interface ethernet1/1/8
 lacp port-priority 4096
 lacp rate fast
 no shutdown
!
...
```

Interfaces

Create a LAG and then add LAG member interfaces. By default, all interfaces are in `no shutdown` and `switchport` modes.

1 Create a LAG in CONFIGURATION mode.

```
interface port-channel port-channel number
```

2 Enter INTERFACE mode.

```
interface ethernet node/slot/port[:subport]
```

3 Set the channel group mode to Active in INTERFACE mode.

```
channel-group number mode active
```

Configure dynamic LAG interfaces

```
OS10(config)# interface port-channel 10
OS10(conf-if-po-10)# exit
OS10(config)# interface ethernet 1/1/10
OS10(conf-if-eth1/1/10)# no switchport
OS10(conf-if-eth1/1/10)# channel-group 10 mode active
OS10(conf-if-eth1/1/10)# exit
OS10(config)# interface ethernet 1/1/11
OS10(conf-if-eth1/1/11)# no switchport
OS10(conf-if-eth1/1/11)# channel-group 10 mode active
```

Rates

Protocol data units (PDUs) are exchanged between port-channel (LAG) interfaces to maintain LACP sessions. PDUs are transmitted at either a slow or fast transmission rate, depending on the LACP timeout value. The timeout value is the amount of time that a LAG interface waits for a PDU from the remote system before bringing the LACP session down.

By default, the LACP rate is `normal` (long timeout). If you configure a `fast` LACP rate, a short timeout sets.

- Set the LACP rate in CONFIGURATION mode.

  ```
  lacp rate [fast | normal]
  ```

Configure LACP timeout

```
OS10(conf-if-eth1/1/29)# lacp rate fast
```

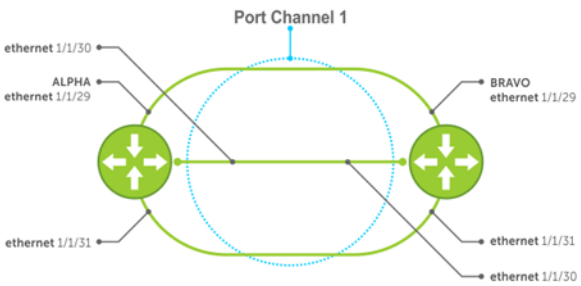
View port status

```
OS10# show lacp port-channel

Port-channel 20 admin up, oper up, mode lacp
Actor System ID: Priority 32768, Address f8:b1:56:00:02:33
Partner System ID: Priority 4096, Address 10:11:22:22:33:33
Actor Admin Key 20, Oper Key 20, Partner Oper Key 10
LACP LAG ID 20 is an aggregatable link
A - Active LACP, B - Passive LACP, C - Short Timeout, D - Long Timeout
E - Aggregatable Link, F - Individual Link, G - IN_SYNC, H - OUT_OF_SYNC,
I - Collection enabled, J - Collection disabled, K - Distribution enabled,
L - Distribution disabled, M - Partner Defaulted, N - Partner Non-defaulted,
O - Receiver is in expired state, P - Receiver is not in expired state
Port ethernet1/1/14 is Enabled, LACP is enabled and mode is lacp
 Actor Admin: State BCFHJKNO Key 20 Priority 32768
     Oper: State BDEGIKNO Key 20 Priority 32768
 Partner Admin: State BCEGIKNP Key 0 Priority 0
     Oper: State BDEGIKNO Key 10 Priority 32768
Port ethernet1/1/16 is Enabled, LACP is enabled and mode is lacp
 Actor Admin: State BCFHJKNO Key 20 Priority 32768
     Oper: State BDEGIKNO Key 20 Priority 32768
 Partner Admin: State BCEGIKNP Key 0 Priority 0
     Oper: State BDEGIKNO Key 10 Priority 32768
```

Sample configuration

This sample topology is based on two routers — Alpha and Bravo.



Alpha LAG configuration summary

```
OS10(config)# interface port-channel 1
OS10(conf-if-po-1)# exit
OS10(config)# interface ethernet 1/1/49
OS10(conf-if-eth1/1/49)# no switchport
OS10(conf-if-eth1/1/49)# channel-group 1 mode active
OS10(conf-if-eth1/1/49)# interface ethernet 1/1/50
OS10(conf-if-eth1/1/50)# no switchport
OS10(conf-if-eth1/1/50)# channel-group 1 mode active
OS10(conf-if-eth1/1/50)# interface ethernet 1/1/51
OS10(conf-if-eth1/1/51)# no switchport
OS10(conf-if-eth1/1/51)# channel-group 1 mode active
```

Bravo LAG configuration summary

```
OS10(config)# interface port-channel 1
OS10(conf-if-po-1)# exit
OS10(config)# interface ethernet 1/1/49
OS10(conf-if-eth1/1/49)# no switchport
OS10(conf-if-eth1/1/49)# channel-group 1 mode active
OS10(conf-if-eth1/1/49)# interface ethernet 1/1/50
OS10(conf-if-eth1/1/50)# no switchport
OS10(conf-if-eth1/1/50)# channel-group 1 mode active
OS10(conf-if-eth1/1/50)# interface ethernet 1/1/51
OS10(conf-if-eth1/1/51)# no switchport
OS10(conf-if-eth1/1/51)# channel-group 1 mode active
```

Alpha verify LAG port configuration

```
OS10# show lacp port-channel

Port-channel 1 admin up, oper up, mode lacp
Actor System ID: Priority 32768, Address 34:17:eb:f2:c7:c4
Partner System ID: Priority 32768, Address 34:17:eb:f2:9b:c4
Actor Admin Key 1, Oper Key 1, Partner Oper Key 1
LACP LAG ID 1 is an aggregatable link
A - Active LACP, B - Passive LACP, C - Short Timeout, D - Long Timeout
E - Aggregatable Link, F - Individual Link, G - IN_SYNC, H - OUT_OF_SYNC,
I - Collection enabled, J - Collection disabled, K - Distribution enabled,
L - Distribution disabled, M - Partner Defaulted, N - Partner Non-defaulted,
O - Receiver is in expired state, P - Receiver is not in expired state
Port ethernet1/1/49 is Enabled, LACP is enabled and mode is lacp
Actor Admin: State BCFHJKNO Key 1 Priority 32768
     Oper: State BDEGIKNO Key 1 Priority 32768
Partner Admin: State BCEGIKNP Key 0 Priority 0
     Oper: State BDEGIKNO Key 1 Priority 32768
Port ethernet1/1/50 is Enabled, LACP is enabled and mode is lacp
Actor Admin: State BCFHJKNO Key 1 Priority 32768
     Oper: State BDEGIKNO Key 1 Priority 32768
Partner Admin: State BCEGIKNP Key 0 Priority 0
     Oper: State BDEGIKNO Key 1 Priority 32768
Port ethernet1/1/51 is Enabled, LACP is enabled and mode is lacp
Actor Admin: State BCFHJKNO Key 1 Priority 32768
     Oper: State BDEGIKNO Key 1 Priority 32768
Partner Admin: State BCEGIKNP Key 0 Priority 0
     Oper: State BDEGIKNO Key 1 Priority 32768
```

Bravo verify LAG port configuration

```
bravo# show interface ethernet 1/1/29

Ethernet 1/1/29 is up, line protocol is up
Port is part of Port-channel
Hardware is Dell EMC Eth, address is 90:b1:1c:f4:9b:a2
    Current address is 90:b1:1c:f4:9b:a2
Pluggable media present, QSFP+ type is QSFP+ 40GBASE CR 1.0M
    Wavelength is 25
    SFP receive power reading is 0.0
```

```
Interface index is 16866812
Internet address is not set
Mode of IPv4 Address Assignment : not set
MTU 1532 bytes, IP MTU  bytes
LineSpeed auto
Flowcontrol rx  tx
ARP type: ARPA, ARP Timeout: 240
Last clearing of show "interface" counters :
Queuing strategy : fifo
Input statistics:
     466 packets, 45298 octets
     224 64-byte pkts,1 over 64-byte pkts, 241 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     466 Multicasts, 0 Broadcasts
     0 runts, 0 giants, 0 throttles
     0 CRC, 0 overrun, 465 discarded
Output statistics:
     7840 packets, 938965 octets
     0 64-byte pkts,1396 over 64-byte pkts, 6444 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     7840 Multicasts, 0 Broadcasts,0 Unicasts
     0 throttles, 0 discarded, 0 Collisions, 0 wreddrops
Rate Info(interval 299 seconds):
     Input 0 Mbits/sec, 0 packets/sec, 0%  of line rate
     Output 0 Mbits/sec, 1 packets/sec, 0%  of line rate
Time since last interface status change : 01:25:29
```

Verify LAG 1

```
OS10# show interface port-channel 1

Port-channel 1 is up,line protocol is  up
Hardware address is  Current address is
Interface index is  85886081
Minimum number of links to bring Port-channel up is 1
Internet address is not set
Mode of IPv4 Address Assignment : not set
Lag MTU is 1500 ,IP MTU  bytes
Linespeed AUTO
Members in this channel ethernet1/1/29 ethernet1/1/30 ethernet1/1/31
ARP type: ARPA    Arp timeout: 240
Last clearing of "show interface" counters :
Queuing strategy :fifo
Input statistics:
1388 packets, 135026 octets
666 64-byte pkts,1 over 64-byte pkts, 721 over 127-byte pkts
0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
1388 Multicasts, 0 Broadcasts
0 runts, 0 giants, 0 throttles
0 CRC, 0 overrun, 1387 discarded
Output statistics:
2121444503 packets, 135773749275 octets
2121421152 64-byte pkts,4182 over 64-byte pkts, 19169 over 127-byte pkts
0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
23351 Multicasts, 0 Broadcasts,2121421152 Unicasts
0 throttles, 143426 discarded, 0 Collisions, 0 wreddrops
Rate Info(interval 299 seconds):
Input 0 Mbits/sec, 0 packets/sec, 0%  of line rate
Output 0 Mbits/sec, 3 packets/sec, 0%  of line rate
Time since last interface status change : 01:24:43
```

Verify LAG status

```
OS10# show lacp port-channel

Port-channel 1 admin up, oper up, mode lacp
Actor System ID: Priority 32768, Address 90:b1:1c:f4:9b:8a
Partner System ID: Priority 32768, Address 00:01:e8:8a:fd:9e
Actor Admin Key 1, Oper Key 1, Partner Oper Key 1
```

```
LACP LAG ID 1 is an aggregatable link

A - Active LACP, B - Passive LACP, C - Short Timeout, D - Long Timeout
E - Aggregatable Link, F - Individual Link, G - IN_SYNC, H - OUT_OF_SYNC,
I - Collection enabled, J - Collection disabled, K - Distribution enabled,
L - Distribution disabled, M - Partner Defaulted, N - Partner Non-defaulted,
O - Receiver is in expired state, P - Receiver is not in expired state

Port ethernet1/1/29 is Enabled, LACP is enabled and mode is lacp
 Actor Admin: State  Key 1 Priority 32768
     Oper: State  Key 1 Priority 32768
 Partner Admin: State  Key 0 Priority 0
     Oper: State  Key 1 Priority 32768
Port ethernet1/1/30 is Enabled, LACP is enabled and mode is lacp
 Actor Admin: State  Key 1 Priority 32768
     Oper: State  Key 1 Priority 32768
 Partner Admin: State  Key 0 Priority 0
     Oper: State  Key 1 Priority 32768
Port ethernet1/1/31 is Enabled, LACP is enabled and mode is lacp
 Actor Admin: State  Key 1 Priority 32768
     Oper: State  Key 1 Priority 32768
 Partner Admin: State  Key 0 Priority 0
     Oper: State  Key 1 Priority 32768
```

Verify LAG membership

```
OS10# show lacp interface ethernet 1/1/29

Interface ethernet1/1/29 is up
  Channel group is 1 port channel is po1
  PDUS sent: 17
  PDUS rcvd: 11
  Marker sent: 0
  Marker rcvd: 0
  Marker response sent: 0
  Marker response rcvd: 0
  Unknown packetse rcvd: 0
  Illegal packetse rcvd: 0
Local Port:     MAC Address=74:e6:e2:f5:b5:80
  System Identifier=32768,32768
  Port Identifier=32768,32768
  Operational key=1
  LACP_Activity=passive
  LACP_Timeout=Long Timeout(30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
  Partner information refresh timeout=Long Timeout(90s)
Actor Admin State=BCFHJKNO
Actor Oper State=BDEGIKNO
Neighbor: 276
  MAC Address=00:00:00:00:00:00
  System Identifier=,00:00:00:00:00:00
  Port Identifier=0,14:18:77:7a:2d:00
  Operational key=1
  LACP_Activity=passive
  LACP_Timeout=Long Timeout(30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
Partner Admin State=BCEGIKNP
Partner Oper State=BDEGIKNO
```

LACP fallback

LACP fallback allows downstream devices, like Servers, connected to ports of a switch configured as LACP to establish a link when the system is not able to finalize the LACP handshake.

For example, when servers boot in PXE mode, the server cannot exchange LACP PDUs and the switch does not enable the ports.

Whenever a PXE server reboots, both the port-channel and ports go down. While rebooting, the ports come up, but not the port-channel. LACP fallback enables the port-channel to be up and keeps sending packets to the PXE server.

When you enable LACP fallback, the switch starts a timer. If the timer expires before LACP completes, then the switch selects one port of the port group and makes it operational.

You can set the timer using the `lacp fallback timeout` *`timer-value`* command.

The LACP fallback feature adds a member port to LACP port-channel if it does not receive LACP PDUs from the peer for a particular period of time.

The server uses the fallback port to finalize the PXE-boot process. When the server starts with the OS, the process completes the LACP handshake and the fallback port re-unites the other members. The member port becomes active and sends packets to the PXE server.

When the switch starts receiving LACP PDU, OS10 ungroups the statically added member port from LACP port-channel and resumes with normal LACP functionality..

When you enable LACP fallback, the port that comes up is selected based on the following:

- LACP port priority configuration allows deterministic port allocation. The port with the least priority is placed in the active state when a port-channel is in LACP fallback mode.
- If all the ports in a port-channel have same port priority, the switch internally compares the interface names by base name, module number, port number, and then selects the lowest one to be active. For example, Ethernet 1 is less than Ethernet 2 and hence Ethernet 1 becomes active.
- In a VLT network, if the interface name is the same on both the VLT peers, then the port in switch with lower system MAC address becomes active.

Limitations

- OS10 switches cannot be a PXE client irrespective of whether it acts as a VLT peer or ToR switch.
- If you are configuring LACP fallback in a VLT domain, configure `lacp fallback` commands in both the VLT peers.
- If you do not enable LACP fallback in one of the VLT peers, or configure different time-out values in the peers, then the switch might behave differently.
- The LACP fallback feature adds or groups a member port to the port channel only when the switch does not receive LACP PDUs from the peer, to make the link connected to the PXE client device as operational. As PXE clients handle untagged DHCP request, you need to configure the LACP fallback only on an untagged VLAN to reach the DHCP/PXE server.
- After the LACP fallback election, if a port with lower priority port is configured to be part of the same port-channel, it would trigger re-election.

Configure LACP fallback

1. Enable LACP fallback with the `lacp fallback enable` command in port-channel INTERFACE mode.
2. Set a timer for receiving LACP PDUs using `lacp fallback timeout` *`timer-value`* in port-channel INTERFACE mode.
3. (Optional) Enable or disable LACP fallback port preemption using `lacp fallback preemption {enable | disable}` in port-channel INTERFACE mode.

Example configuration

```
OS10# configure terminal
OS10(config)# interface port-channel 1
```

```
OS10(conf-if-po-1)# lacp fallback enable
OS10(conf-if-po-1)# lacp fallback timeout 20
OS10(conf-if-po-1)# lacp fallback preemption enable
```

View LACP fallback configuration

```
OS10# show port-channel summary

Flags:  D - Down    I - member up but inactive    P - member up and active

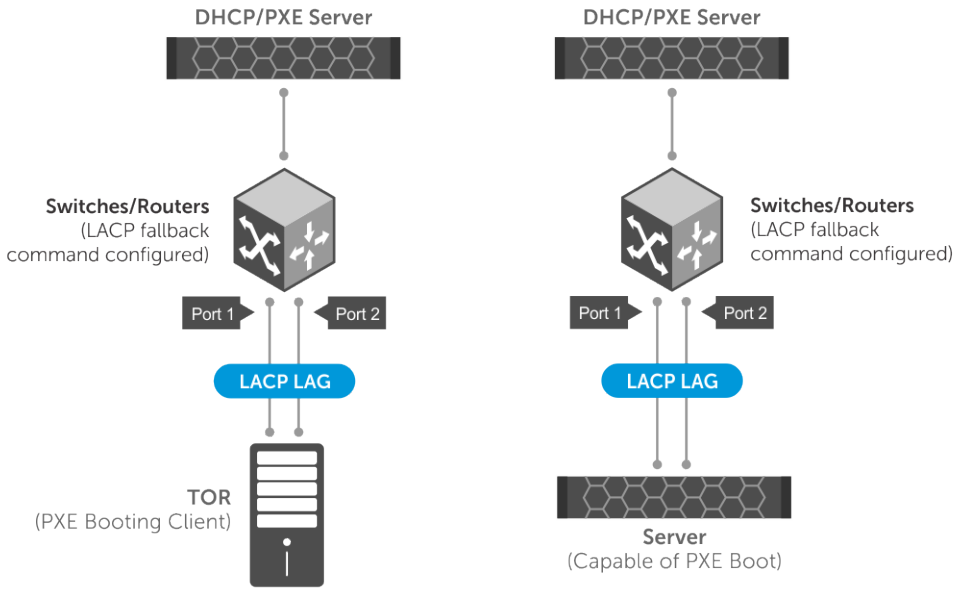
        U - Up (port-channel) F - Fallback enabled

--------------------------------------------------------------------------------
Group Port-Channel      Type     Protocol  Member Ports
--------------------------------------------------------------------------------

1    port-channel1   (UF) Eth      DYNAMIC   1/1/10(P) 1/1/11(I)
```

LACP fallback in non-VLT network

In a non-VLT network, LACP fallback enables rebooting of ToR or server connected to the switch through normal LACP. The other end of the switch is connected to a DHCP/PXE server, as shown in the following illustration:

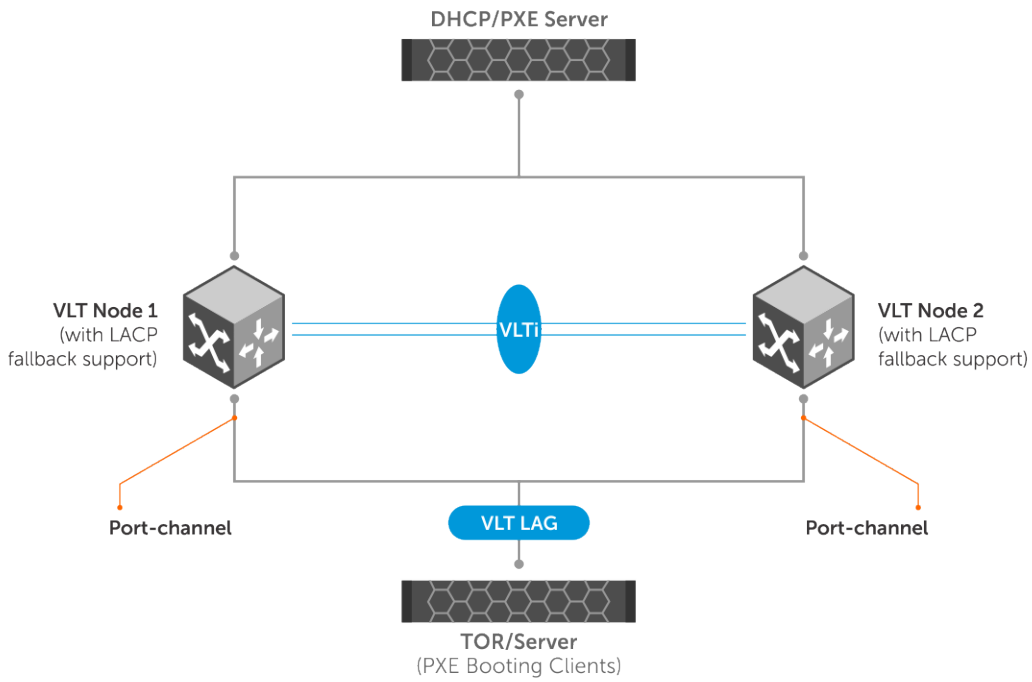


In the above scenario, LACP fallback works as follows:

1. The ToR/server boots up.
2. The switch detects the link that is up and checks fallback enabled status. If fallback is enabled, the device waits for the time-out period for any LACP BPDUs. If there are no LACP BPDUs received within the time period, then the LAG enters into fallback mode and adds the first operationally UP port to the port-channel instead of placing it in an inactive state.
3. Now the ToR/server has one port up and active. The active port sends packets to the DHCP/PXE server.
4. After receiving response from the DHCP server, the ToR/server proceeds to boot from the TFTP/NFS server.
5. When the ToR/server is fully loaded with the boot image and configurations, the server starts sending LACP PDUs.
6. When the switch receives LACP PDUs from ToR/server, the device comes out of the fallback mode and activates the LAG through normal LACP process.

LACP fallback in VLT domain

In a VLT domain, LACP fallback enables rebooting of ToR or server connected to VLT nodes through VLT port-channel. The other end of the VLT nodes are connected to a DHCP/PXE server, as shown in the following illustration:



In the above scenario, LACP fallback works as follows:

1 The ToR/server boots up.
2 One of the VLT peers takes care of controlling the LACP fallback mode. All events are sent to the controlling VLT peer for deciding the port that should be brought up and then the decision is passed on to peer devices.
3 The controlling VLT peer can decide to bring up one of the ports in either the local port-channel or in the peer VLT port-channel.
4 One of the ports, local or peer, becomes active based on the decision of the controlling VLT peer.
5 Now the ToR/server has one port up and active. The active port sends packets to the DHCP/PXE server.
6 After receiving response from the DHCP server, the ToR/server proceeds to boot from the TFTP/NFS server.
7 When the ToR/server is fully loaded with the boot image and configurations, the server starts sending LACP PDUs.
8 When the switch receives LACP PDUs from ToR/server, the controlling VLT peer makes the LACP port to come out of the fallback mode and to resume the normal functionality.

LACP commands

channel-group

Assigns and configures a physical interface to a port-channel group.

Syntax `channel-group` *number* `mode {active | on | passive}`

Parameters

- *number* — Enter the port-channel group number (1 to 128). The maximum number of port-channels is 128. The maximum physical port/maximum NPU is supported.
- mode — Enter the interface port-channel mode.
- active — Enter to enable the LACP interface. The interface is in the Active Negotiating state when the port starts negotiations with other ports by sending LACP packets.
- on — Enter so that the interface is not part of a dynamic LAG but acts as a static LAG member.
- passive — Enter to only enable LACP if it detects a device. The interface is in the Passive Negotiation state when the port responds to the LACP packets that it receives but does not initiate negotiation until it detects a device.

Default Not configured

Command Mode INTERFACE

Usage Information When you delete the last physical interface from a port-channel, the port-channel remains. Configure these attributes on an individual member port. If you configure a member port with an incompatible attribute, OS10 suspends that port in the port-channel. The member ports in a port-channel must have the same setting for link speed capability and duplex capability. The no version of this command removes the interface from the port-channel.

Example

```
OS10(config)# interface ethernet 1/1/10
OS10(conf-if-eth1/1/10)# channel-group 10 mode active
OS10(conf-if-eth1/1/10)# exit
OS10(config)# interface ethernet 1/1/11
OS10(conf-if-eth1/1/11)# channel-group 10 mode active
```

Supported Releases 10.2.0E or later

clear lacp counters

Clears the statistics for all interfaces for LACP groups.

Syntax `clear lacp counters [interface port-channel channel-number]`

Parameters

- interface port-channel — (Optional) Enter the interface port-channel number.
- *channel-number* — (Optional) Enter the LACP port-channel number (1 to 128).

Default Not configured

Command Mode EXEC

Usage Information If you use this command for a static port-channel group without enabling the aggregation protocol, the device ignores the command. If you do not enter a port-channel number, the LACP counters for all LACP port groups clear.

Example

```
OS10# clear lacp counters
```

Example (Port-Channel)

```
OS10# clear lacp counters interface port-channel 20
```

Supported Releases 10.2.0E or later

lacp fallback enable

Enables LACP fallback mode.

Syntax `lacp fallback enable`

Parameters None

Default Disabled

Command Mode Port-channel INTERFACE

Usage Information The `no` version of this command disables LACP fallback mode.

Example

```
OS10# configure terminal
OS10(config)# interface port-channel 1
OS10(conf-if-po-1)# lacp fallback enable
```

Supported Releases 10.3.2E(R3) or later

lacp fallback preemption

Enables or disables LACP fallback port preemption.

Syntax `lacp fallback preemption {enable | disable}`

Parameters

- `enable`—Enables preemption on the port-channel.
- `disable`—Disables preemption on the port-channel.

Default Enabled

Command Mode Port-channel INTERFACE

Usage Information When you enable preemption, the fallback port election preempts the already elected fallback port and elects a new fallback port.

The new port is elected based on the following events:

- When a non-fallback port configured with low priority.
- When a low-priority port becomes operationally UP.
- When a port with the least numbering is operationally UP.
- A port with the lowest priority is elected as fallback port, if non-default LACP port priority is configured on a port even though preemption is disabled.
- The `lacp fallback preemption disable` command is not applicable on port priority events that you have configured or triggered.

Example

```
OS10# configure terminal
OS10(config)# interface port-channel 1
OS10(conf-if-po-1)# lacp fallback preemption enable
```

```
OS10# configure terminal
OS10(config)# interface port-channel 1
OS10(conf-if-po-1)# lacp fallback preemption disable
```

Supported Releases 10.4.1.0 or later

lacp fallback timeout

Configures LACP fallback time out period.

Syntax	`lacp fallback timeout` *timer-value*
Parameters	*timer-value*—Enter the timer values in seconds, ranging from 0 to 100 seconds.
Default	15 seconds
Command Mode	Port-channel INTERFACE
Usage Information	The `no` version of this command returns the timer to default value.

Example	`OS10# configure terminal`
	`OS10(config)# interface port-channel 1`
	`OS10(conf-if-po-1)# lacp fallback timeout 20`

Supported Releases	10.3.2E(R3) or later

lacp max-bundle

Configures the maximum number of active members allowed in a port-channel.

Syntax	`lacp max-bundle` *max-bundle-number*
Parameters	*max-bundle-number* — Enter the maximum bundle size (1 to 32).
Default	32
Command Mode	INTERFACE
Usage Information	The `no` version of this command resets the maximum bundle size to the default value.
Example	`OS10(conf-if-po-10)# lacp max-bundle 10`
Supported Releases	10.2.0E or later

lacp port-priority

Sets the priority for the physical interfaces for LACP.

Syntax	`lacp port-priority` *priority*
Parameters	*priority* — Enter the priority for the physical interfaces (0 to 65535).
Default	32768
Command Mode	INTERFACE
Usage Information	LACP uses the port priority with the port number to create the port identifier. The port priority decides which ports are put into Standby mode when there is a hardware limitation that prevents all compatible ports from aggregating, or when you have more than eight ports configured for the channel group. When setting the priority, a higher number means a lower priority. The `no` version of this command returns the port priority to the default value.
Example	`OS10(conf-range-eth1/1/7-1/1/8)# lacp port-priority 32768`

Supported Releases	10.2.0E or later

lacp rate

Sets the rate at which LACP sends control packets.

Syntax	`lacp rate {fast | normal}`
Parameters	• `fast` — Enter the fast rate of 1 second. • `normal` — Enter the default rate of 30 seconds.
Default	30 seconds
Command Mode	INTERFACE
Usage Information	Change the LACP timer rate to modify the duration of the LACP timeout. The `no` version of this command resets the rate to the default value.
Example	`OS10(conf-range-eth1/1/7-1/1/8)# lacp rate fast`
Supported Releases	10.2.0E or later

lacp system-priority

Sets the system priority of the device for LACP.

Parameters	*priority* — Enter the priority value for physical interfaces (0 to 65535).
Default	32768
Command Mode	CONFIGURATION
Usage Information	Each device that runs LACP has an LACP system priority value. LACP uses the system priority with the MAC address to form the system ID and also during negotiation with other systems. The system ID is unique for each device. The `no` version of this command resets the system priority to the default value.
Example	`OS10(config)# lacp system-priority 32768`
Supported Releases	10.2.0E or later

show lacp counter

Displays information about LACP statistics.

Syntax	`show lacp counter [interface port-channel` *channel-number*`]`
Parameters	• `interface port-channel` — (Optional) Enter the interface port-channel. • *channel-number* — (Optional) Enter the LACP channel group number (1 to 128).
Default	Not configured
Command Mode	EXEC
Usage Information	All channel groups display if you do not enter the *channel-number* parameter.

Example

```
OS10# show lacp counter interface port-channel 1

LACPDUs        Marker         Marker Response   LACPDUs
Port           Sent    Recv   Sent    Recv  Sent   Recv  Pkts Err
--------------------------------------------------------------------
port-channel1
Ethernet1/1    554      536   0       0     0      0      0
Ethernet1/2    527      514   0       0     0      0      0
Ethernet1/3    535      520   0       0     0      0      0
Ethernet1/4    515      502   0       0     0      0      0
Ethernet1/5    518      505   0       0     0      0      0
Ethernet1/6    540      529   0       0     0      0      0
Ethernet1/7    541      530   0       0     0      0      0
Ethernet1/8    547      532   0       0     0      0      0
Ethernet1/9    544      532   0       0     0      0      0
Ethernet1/10   513      501   0       0     0      0      0
Ethernet1/11   497      485   0       0     0      0      0
Ethernet1/12   493      486   0       0     0      0      0
Ethernet1/13   492      485   0       0     0      0      0
--more--
```

Supported Releases 10.2.0E or later

show lacp interface

Displays information about specific LACP interfaces.

Syntax `show lacp interface ethernet` *node/slot/port*

Parameters *node/slot/port* — Enter the interface information.

Default Not configured

Command Mode EXEC

Usage Information The *LACP_activity* field displays if you configure the link in Active or Passive port-channel mode. The *Port Identifier* field displays the port priority as part of the information including the port number. For example, `Port Identifier=0x8000,0x101`, where the port priority value is `0x8000` and the port number value is `0x101`.

Example

```
OS10# show lacp interface ethernet 1/1/129
Invalid Port id, Max. Port Id is: 32
OS10# show lacp interface ethernet 1/1/29

Interface ethernet1/1/29 is up
  Channel group is 1 port-channel is po1
  PDUS sent: 365
  PDUS rcvd: 17
  Marker sent: 0
  Marker rcvd: 0
  Marker response sent: 0
  Marker response rcvd: 0
  Unknown packetse rcvd: 0
  Illegal packetse rcvd: 0
Local Port: ethernet1/1/29    MAC Address=90:b1:1c:f4:9b:8a
  System Identifier=32768,32768
  Port Identifier=32768,32768
  Operational key=1
  LACP_Activity=passive
  LACP_Timeout=Long Timeout(30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
  Partner information refresh timeout=Long Timeout(90s)
Actor Admin State=BCFHJKNO
Actor Oper State=BDEGIKNO
```

```
Neighbor: 178
  MAC Address=00:00:00:00:00:00
  System Identifier=,00:00:00:00:00:00
  Port Identifier=0,00:01:e8:8a:fd:9e
  Operational key=1
  LACP_Activity=passive
  LACP_Timeout=Long Timeout(30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
Partner Admin State=BCEGIKNP
Partner Oper State=BDEGIKMO
```

Supported Releases 10.2.0E or later

show lacp neighbor

Displays information about LACP neighbors.

Syntax `show lacp neighbor [interface port-channel channel-number]`

Parameters

- `interface port-channel` — (Optional) Enter the interface port-channel.
- *`channel-number`* — (Optional) Enter the port-channel number for the LACP neighbor (1 to 128).

Default Not configured

Command Mode EXEC

Usage Information All channel groups display if you do not enter the *`channel-number`* parameter.

Example

```
OS10# show lacp neighbor interface port-channel 1

Flags:S-Device is sending Slow LACPDUs F-Device is sending Fast LACPdus
      A-Device is in Active mode       P-Device is in Passive mode
Port-channel port-channel1 neighbors
Port: ethernet1/1/29
Partner System Priority: 32768
Partner System ID: 00:01:e8:8a:fd:9e
Partner Port: 178
Partner Port Priority: 32768
Partner Oper Key: 1
Partner Oper State:aggregation synchronization collecting distributing
defaulted expired
```

Supported Releases 10.2.0E or later

show lacp port-channel

Displays information about LACP port-channels.

Syntax `show lacp port-channel [interface port-channel channel-number]`

Parameters

- `interface port-channel` — (Optional) Enter the interface port-channel.
- *`channel-number`* — (Optional) Enter the port-channel number for the LACP neighbor (1 to 128).

Default Not configured

Command Mode EXEC

Usage Information All channel groups display if you do not enter the *channel-number* parameter.

Example

```
OS10# show lacp port-channel 1

Port-channel 1 admin up, oper up, mode lacp
Actor System ID: Priority 32768, Address 90:b1:1c:f4:9b:8a
Partner System ID: Priority 32768, Address 00:01:e8:8a:fd:9e
Actor Admin Key 1, Oper Key 1, Partner Oper Key 1
LACP LAG ID 1 is an aggregatable link
A-Active LACP, B-Passive LACP, C-Short Timeout, D-Long Timeout
E-Aggregatable Link, F-Individual Link, G-IN_SYNC, H-OUT_OF_SYNC,
I-Collection enabled, J-Collection disabled, K-Distribution enabled,
L-Distribution disabled, M-Partner Defaulted, N-Partner Non-defaulted,
O-Receiver is in expired state, P-Receiver is not in expired state
Port ethernet1/1/29 is Enabled, LACP is enabled and mode is lacp
 Actor Admin: State BCFHJKNO Key 1 Priority 32768
     Oper: State BDEGIKNO Key 1 Priority 32768
 Partner Admin: State BCEGIKNP Key 0 Priority 0
     Oper: State BDEGIKMO Key 1 Priority 32768
```

Supported Releases 10.2.0E or later

show lacp system-identifier

Displays the LACP system identifier for a device.

Syntax `show lacp system-identifier`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information The LACP system ID is a combination of the configurable LACP system priority value and the MAC address. Each system that runs LACP has an LACP system priority value. The default value is 32768 or configure a value between 1 and 65535. LACP uses the system priority with the MAC address to form the system ID and uses the system priority during negotiation with other devices. A higher system priority value means a lower priority. The system ID is different for each device.

Example

```
OS10# show lacp system-identifier

Actor System ID: Priority 32768, Address 90:b1:1c:f4:9b:8a
```

Supported Releases 10.2.0E or later

Link Layer Discovery Protocol

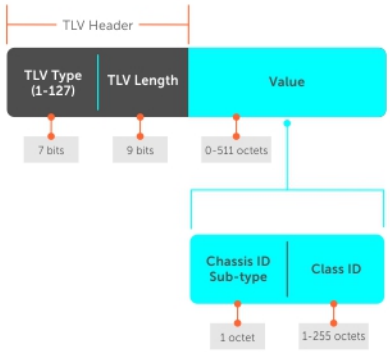
Link layer discovery protocol (LLDP) enables a local area network (LAN) device to advertise its system and receive system information from adjacent LAN devices.

- LLDP is enabled by default on OS10 interfaces.
- An LLDP-enabled interface supports up to eight neighbors. An OS10 switch supports a maximum of 250 neighbors per system.
- OS10 devices receive and periodically transmit Link Layer Discovery Protocol Data Units (LLDPDUs), which are data packets. The default transmission interval is 30 seconds.
- LLDPDU information received from a neighbor expires after the default time to live (TTL) value of 120 seconds.
- Spanning-tree *blocked* ports allow LLDPDUs.
- 802.1X-controlled ports do not allow LLDPDUs until the connected device is authenticated.
- Link layer discovery protocol-media endpoint discovery (LLDP-MED) is enabled on all interfaces by default.

Protocol data units

LLDP devices exchange system information represented as type, length, and value (TLV) segments:

Type	Information included in the TLV.
Length	Value in bytes of the TLV after the Length field.
Value	System information the agent advertises.



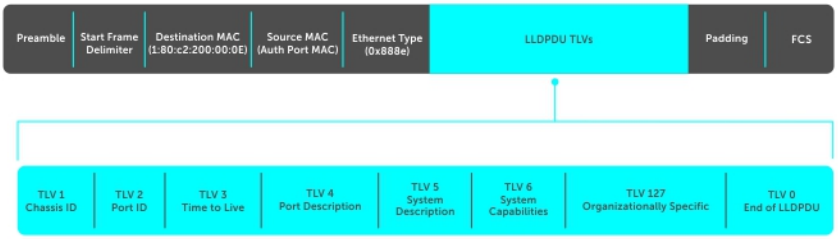
tlv segment

LAN devices transmit LLDPDUs, which encapsulate TLVs, to neighboring LAN devices. LLDP is a one-way protocol and LAN devices (LLDP agents) transmit and/or receive advertisements but they cannot solicit and do not respond to advertisements.

There are three mandatory TLVs followed by zero or more optional TLVs and the end of the LLDPDU TLV. The three mandatory TLVs must be located at the beginning of the LLDPDU in the following order:

- Chassis ID TLV
- Port ID TLV
- Time-to-live TLV

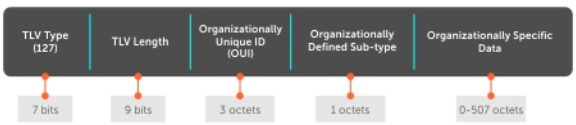
0 — End of LLDPDU	Marks the end of an LLDPDU.
1 — Chassis ID	Identifies the LAN agent.
2 — Port ID	Identifies a port through which the LAN device transmits LLDPDUs.
3 — Time-to-live	Number of seconds that the recipient LLDP agent considers the information associated with this MAP identifier to be valid.
— Optional	Includes sub-types of TLVs that advertise specific configuration information. These sub-types are management TLVs, IEEE 802.1, IEEE 802.3, and TIA-1057 organization-specific TLVs.



Optional TLVs

OS10 supports basic TLVs, IEEE 802.1, and 802.3 organizationally-specific TLVs, and TIA-1057 organizationally-specific TLVs. A basic TLV is an optional TLV sub-type. This kind of TLV contains essential management information about the sender.

A professional organization or vendor defines organizationally-specific TLVs. They have two mandatory fields, in addition to the basic TLV fields.



Organizationally-specific TLVs

Table 5. Basic management TLVs

TLV	Type	Description
Port description	4	User-defined alphanumeric string that describes the port (port ID or interface description).
System name	5	User-defined alphanumeric string that identifies the system.
System description	6	Includes the following information: • Host description • Dell OS version • Dell application software version • Build time stamp
System capabilities	7	Determines the capabilities of the system.
Management address	8	Network address of the management interface.

There are eight TLV types defined by the 802.1 and 802.3 working groups as a basic part of LLDP. Configure OS10 to advertise any or all of these TLVs.

Table 6. 802.1x organizationally-specific TLVs (Type – 127, OUI – 00-80-C2)

TLV	Subtype	Description
Link aggregation	7	Indicates whether the link associated with the port on which the LLDPDU is transmitted is aggregated. Also indicates whether the link is currently aggregated and provides the aggregated port identifier if the link is aggregated.
Port VLAN ID	1	Untagged VLAN to which a port belongs.
Protocol identity	4	Not supported.

Table 7. 802.3 organizationally-specific TLVs (Type – 127, OUI – 00-12-0F)

TLV	Subtype	Description
MAC/PHY configuration/status	1	Indicates duplex and bit rate capability and the current duplex and bit rate settings of the sending device. Also indicates whether the current settings are due to auto-negotiation or manual configuration.
Power via MDI	2	Not supported.
Maximum frame size	4	Maximum frame size capability of the MAC and PHY.

Table 8. iDRAC organizationally-specific TLVs; Subtypes used in iDRAC custom TLVs (Type – 127, OUI – 0xF8-0xB1-0x56)

TLV	Subtype	Description
Originator	1	Indicates the iDRAC string that is used as originator. This string enables external switches to easily identify iDRAC LLDP PDUs.
Port type	2	Following are the applicable port types: 1 iDRAC port (dedicated) 2 NIC port 3 iDRAC and NIC port (shared)
Port FQDD	3	Port number that uniquely identifies a NIC port within a server.
Server service tag	4	Service tag ID of the server.
Server model name	5	Model name of the server. For example, PowerEdge FC640.
Server slot number	6	Slot number of the server. For example, 1, 2, 3, 1a, 1b, and so on.
Chassis service tag	7	Service tag ID of the chassis. (Applicable only to blade servers.)

TLV	Subtype	Description
Chassis model	8	Model name of the chassis. (Applicable only to blade servers.)
IOM service tag	9	Service tag ID of the IOM device. (Applicable only to blade servers.)
IOM model name	10	Model name of the IOM device. (Applicable only to blade servers.)
IOM slot label	11	Slot label of the IOM device. For example, A1, B1, A2, B2, and so on (applicable only to blade servers).
IOM port number	12	Port number of the NIC. For example, 1, 2, 3, and so on.

Table 9. Isilon-related TLVs (Type – 127, OUI – 0xF8-0xB1-0x56)

TLV	Subtype	Description
Subtypes used in LLDP custom TLVs that are transacted by the TLV nodes		
Originator	1	Indicates the Isilon string that is used as originator. This string enables the OS10 switches to easily identify the Isilon originated LLDP PDUs.
RA prefix	2	Indicates the IPV6 address prefix for SLAAC. This prefix is also used by Isilon nodes to communicate with the master and the OS10 switch to compute the Virtual IP address for the specific fabric instance. The RA prefix is different for each fabric.
Fabric ID	3	Indicates the ID of the fabric the LLDP PDU is originating from.
Isilon-related TLVs – Subtypes used in LLDP custom TLVs that are transacted by the OS10 switches		
Originator	1	Indicates the OS10 string that is used as originator. The string enables the OS10 switches to easily identify LLDP PDUs.
Role	2	Following are the applicable roles: 1 LEAF 2 SPINE 3 UNKNOWN
IP address	3	Indicates the IPv6 address of the originator.
Virtual IP address of the fabric	4	Virtual IP address of the master node. The Isilon nodes can also use this IPv6 address when needed.
MAC address of the physical interface	5	MAC address used by the OS10 switches for ND.

Table 10. Service tag TLV (Type – 127, OUI – 0xF8-0xB1-0x56)

TLV	Subtype	Description
Service tag	21	Indicates the service tag associated with the device.

Table 11. Solution ID TLVs (Type – 127, OUI – 0xF8-0xB1-0x56)

TLV	Subtype	Description
Product base	22	Indicates the product base.
Product serial number	23	Indicates the product serial number.
Product part number	24	Indicates the product part number.

Media endpoint discovery

LLDP-MED provides additional organizationally-specific TLVs to allow endpoint devices and network-connectivity devices to advertise their characteristics and configuration information.

LLDP-MED devices are located at the IEEE 802 LAN network edge and participate in IP communication service using the LLDP-MED framework, such as IP phones and conference bridges. LLDP-MED network connectivity devices provide access to the IEEE 802-based LAN infrastructure for LLDP-MED endpoint devices, such as IP phones. An OS10 device acts as an LLDP-MED network connectivity device.

LLDP-MED provides network connectivity devices to:

- Manage inventory
- Manage power over ethernet (PoE)
- Identify physical location
- Identify network policy

NOTE: Only the Rx function is supported for managing PoE and identifying the physical location. LLDP-MED is designed for but not limited to VoIP endpoints.

Network connectivity device

OS10 acts as an LLDP-MED network-connectivity device (Type 4). Network connectivity devices transmit an LLDP-MED capability TLV to endpoint devices and stores information that endpoint devices advertise.

127/1 — LLDP-MED capabilities	• If the transmitting device supports LLDP-MED • What LLDP-MED TLVs are supported • LLDP device class
127/2 — Network policy	Application type, VLAN ID, L2 priority, and DSCP value.
127/3 — Local identification	Physical location of the device expressed in one of three formats: • Coordinate-based LCI

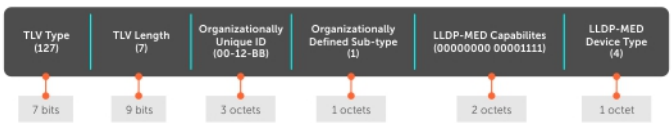
- Civic address LCI
- Emergency call services ELIN

127/4 — Extended power-via-MDI	Power requirements, priority, and power status.

LLDP-MED capabilities TLV

The LLDP-MED capabilities TLV communicates the types of TLVs that the endpoint device and network-connectivity device support. The value of the LLDP-MED capabilities field in the TLV is a 2–octet bitmap. Each bit represents an LLDP-MED capability.

LLDP-MED is enabled by default on an interface. If you disable LLDP-MED, use the `lldp med enable` command to re-enable it on an interface. The device transmits MED PDUs only when it receives a TLV from a peer. The device does not otherwise send PDUs — even if you enable MED on an interface.



LLDP-MED capabilities

Bit 0	LLDP-MED capabilities
Bit 1	Network policy
Bit 2	Location ID
Bit 3	Extended power via MDI-PSE
Bit 4	Extended power via MDI-PD
Bit 5	Inventory
Bits 6-15	Reserved

LLDP-MED device types

0	Type not defined
1	Endpoint class 1
2	Endpoint class 2
3	Endpoint class 3
4	Network connectivity
5-255	Reserved

Network policies TLVs

A network policy in the context of LLDP-MED is a device's VLAN configuration and associated L2 and L3 configurations.

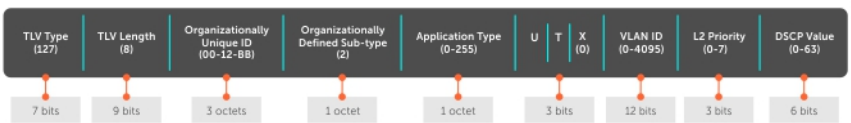
LLDP-MED network policies TLV include:

- VLAN ID
- VLAN tagged or untagged status
- L2 priority
- DSCP value

An integer represents the application type the Type integer shown in the following table, which indicates a device function where a unique network policy is defined. An individual LLDP-MED network policy TLV generates for each application type that you use with OS10 commands, see Advertise LLDP-MED TLVs.

NOTE: Signaling is a series of control packets that exchange between an endpoint device and a network-connectivity device to establish and maintain a connection. These signal packets might require a different network policy than the media packets where a connection is made. In this case, configure the signaling application.

0 — Reserved	—
1 — Voice	Used for dedicated IP telephony handsets and other appliances supporting interactive voice services.
2 — Voice signaling	Used only if voice control packets use a separate network policy than voice data.
3 — Guest voice	Used only for a separate limited voice service for guest users with their own IP telephony handsets and other appliances supporting interactive voice services.
4 — Guest voice signaling	Used only if guest voice control packets use a separate network policy than voice data.
5 — SoftPhone voice	Used for softphone application on a device such as a PC or laptop. This class does not support multiple VLANs and if required, uses an untagged VLAN or a single tagged data-specific VLAN.
6 — Video conferencing	Used only for dedicated video conferencing and similar appliances supporting real-time interactive video.
7 — Streaming video	Used for broadcast- or multicast-based video content distribution and similar applications supporting streaming video services that require specific network policy treatment.
8 — Video signaling	Used only if video control packets use a separate network policy than the video data.
9-255 — Reserved	—



Define network policies

You can manually define LLDP-MED network policies. LLDP commands that you configure in CONFIGURATION mode are global and affect all interfaces. LLDP commands you configure in INTERFACE mode affect only the specific interface.

Create a maximum of 32 network policies and attach the LLDP-MED network policies to a port in CONFIGURATION mode.

- Define the LLDP-MED network policy in CONFIGURATION mode.

```
lldp-med network-policy number app {voice | voice-signaling | guest-voice | guestvoice-
signaling | softphone-voice | streaming-video | video-conferencing | video-signaling}{vlan
vlan-id vlan-type {tag | untag} priority priority dscp dscp value}
```

Configure LLDP-MED network policy for voice applications

```
OS10(config)# lldp med network-policy 10
OS10(config)# lldp med network-policy 10 app
OS10(config)# lldp med network-policy 10 app voice
OS10(config)# lldp med network-policy 1 app voice vlan 10 vlan-type tag
OS10(config)# lldp med network-policy 1 app voice-signaling vlan 10 vlan-type tag priority 2
dscp 1
```

Packet timer values

LLDPDUs transmitt periodically. You can configure LLDP packet timer values for LLPDU transmission.

1 Configure the LLDP packet timer value in CONFIGURATION mode.

```
lldp timer
```

2 Enter the multiplier value for the hold time in CONFIGURATION mode.

```
lldp holdtime-multiplier
```

3 Enter the delay in seconds for LLDP initialization on any interface in CONFIGURATION mode.

```
lldp reinit
```

Configure LLDPDU timer

```
OS10(config)# lldp timer 60
OS10(config)# do show lldp timers
LLDP Timers:
Holdtime in seconds: 120
Reinit-time in seconds: 2
Transmit interval in seconds: 60
```

Configure LLDPDU intervals

```
OS10(config)# lldp holdtime-multiplier 2
OS10(config)# do show lldp timers
LLDP Timers:
Holdtime in seconds: 60
Reinit-time in seconds: 2
Transmit interval in seconds: 30
```

Disable and re-enable LLDP

By default, LLDP is enabled for each interface and globally. You can disable LLDP on an interface or globally. If you disable LLDP globally, LLDP is disabled on all interfaces irrespective of whether LLDP is previously enabled or disabled on an interface. When you enable LLDP globally, the LLDP configuration in interface mode takes precedence over the global LLDP configuration.

1 Disable LLDPDU transmit or receive in INTERFACE mode.

```
no lldp transmit
no lldp receive
```

2 Disable LLDP holdtime multiplier value in CONFIGURATION mode.

```
no lldp holdtime-multiplier
```

3 Disable LLDP initialization in CONFIGURATION mode.

```
no lldp reinit
```

4 Disable LLDP MED in CONFIGURATION or INTERFACE mode.

```
no lldp med
```

5 Disable LLDP TLV in INTERFACE mode.

```
no lldp tlv-select
```

6 Disable LLDP globally in CONFIGURATION mode.

```
no lldp enable
```

Disable LLDP

```
OS10(config)# no lldp timer 100
OS10(config)# no lldp holdtime-multiplier 10
OS10(config)# no lldp reinit 8
```

Disable LLDP interface

```
OS10(config)# interface ethernet 1/1/4
OS10(conf-if-eth1/1/4)# no lldp med
OS10(conf-if-eth1/1/4)# no lldp tlv-select
OS10(conf-if-eth1/1/4)# no lldp transmit
OS10(conf-if-eth1/1/4)# no lldp receive
```

Enable LLDP

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# lldp transmit
OS10(conf-if-eth1/1/1)# lldp receive
```

Disable LLDP globally

```
OS10(config)# no lldp enable
```

Disable and re-enable LLDP on management ports

By default, LLDP is enabled on management ports. You can disable or enable the following LLDP configurations on management ports.

1 Disable the LLDPDU transmit or receive.

```
no lldp transmit
no lldp receive
```

2 Disable LLDP TLVs.

```
no lldp tlv-select basic-tlv {port-description | system-name | system-description | system-
capabilities | management-address}
no lldp tlv-select dot1tlv port-vlan-id
```

Disable LLDP transmit or receive

```
OS10(config)# interface mgmt 1/1/1
OS10(conf-if-ma-1/1/1)# no lldp transmit
OS10(conf-if-ma-1/1/1)# no lldp receive
```

Enable LLDP transmit or receive

```
OS10(config)# interface mgmt 1/1/1
OS10(conf-if-ma-1/1/1)# lldp transmit
OS10(conf-if-ma-1/1/1)# lldp receive
```

Disable LLDP TLVs

```
OS10(config)# interface mgmt 1/1/1
OS10(conf-if-ma-1/1/1)# no lldp tlv-select basic-tlv system-name system-description
OS10(conf-if-ma-1/1/1)# no lldp tlv-select dot1tlv port-vlan-id
```

Enable LLDP TLVs

```
OS10(config)# interface mgmt 1/1/1
OS10(conf-if-ma-1/1/1)# lldp tlv-select basic-tlv system-name system-description
OS10(conf-if-ma-1/1/1)# lldp tlv-select dot1tlv port-vlan-id
```

Advertise TLVs

Configure the system to advertise TLVs from all interfaces or specific interfaces. If you configure an interface, only the interface sends LLDPDUs with the specified TLVs.

1 Enable basic TLVs attributes to transmit and receive LLDP packets in INTERFACE mode.

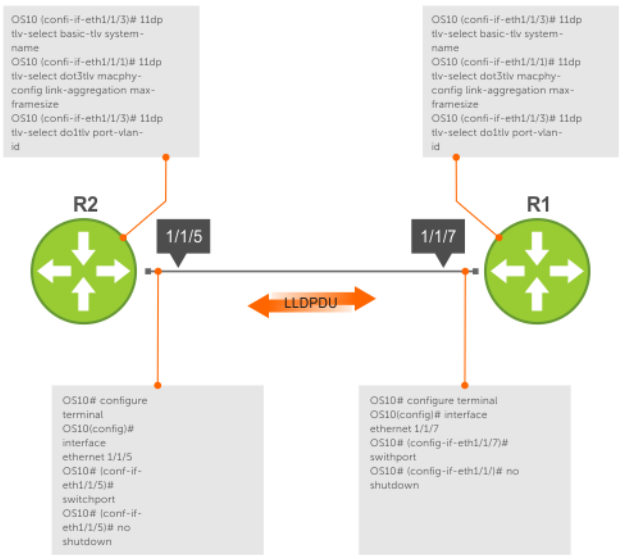
```
lldp tlv-select basic-tlv {port-description | system-name | system-description | system-capabilities | management-address}
```

2 Enable dot3 TLVs to transmit and receive LLDP packets in INTERFACE mode.

```
lldp tlv-select dot3tlv {macphy-config | max-framesize}
```

3 Enable dot1 TLVs to transmit and receive LLDP packets in INTERFACE mode.

```
lldp tlv-select dot1tlv { port-vlan-id | link-aggregation}
```



Configure advertise TLVs

```
OS10(conf-if-eth1/1/3)# lldp tlv-select basic-tlv system-name
OS10(conf-if-eth1/1/1)# lldp tlv-select dot3tlv macphy-config max-framesize
OS10(conf-if-eth1/1/3)# lldp tlv-select dot1tlv link-aggregation
```

Network policy advertisement

LLDP-MED is enabled on all interfaces by default. Configure OS10 to advertise LLDP-MED TLVs from configured interfaces. Define LLDP-MED network policies before applying the policies to an interface. Attach only one network policy per interface.

- Define an LLDP-MED network-policy on an interface in CONFIGURATION mode.

  ```
  lldp-med network-policy {add | remove} number
  ```

  - `add` — Attach the network policy to an interface.
  - `remove` — Remove the network policy from an interface.
  - *`number`* — Enter a network policy index number, from 1 to 32.

Configure advertise LLDP-MED network policies

```
OS10(conf-if-eth1/1/5)# lldp-med network-policy add 1
```

Fast start repeat count

Fast start repeat count enables a network-connectivity device to advertise itself at a faster rate for a limited amount of time. The fast start timer starts when a network-connectivity device receives the first LLDP frame from a newly detected endpoint.

When an LLDP-MED endpoint is newly detected or connected to the network, the `lldp-med fast-start-repeat-count` command enables the network to quickly detect the endpoint. The LLDP-MED fast start repeat count specifies the number of LLDP packets that send during the LLDP-MED fast start period. By default, the device sends three packets per interval. Change the number of packets a device sends per second a maximum of 10.

Rapid availability is crucial for applications such as emergency call service location (E911).

- Enable fast start repeat count which is the number of packets sent during activation in CONFIGURATION mode, from 1 to 10, default 3.
  `lldp-med fast-start-repeat-count` *number*

Configure fast start repeat count

```
OS10(config)# lldp med fast-start-repeat-count 5
```

View LLDP configuration

- View the LLDP configuration in EXEC mode.
  `show running-configuration`
- View LLDP error messages in EXEC mode.
  `show lldp errors`
- View LLDP timers in EXEC mode.
  `show lldp timers`
- View the LLDP traffic in EXEC mode.
  `show lldp traffic`

View running configuration

```
OS10# show running-configuration
```

View LLDP errors

```
OS10# show lldp errors
Total Memory Allocation Failures : 0
Total Input Queue Overflows : 0
Total Table Overflows : 0
```

View LLDP timers

```
OS10# show lldp timers
LLDP Timers:
Holdtime in seconds: 120
Reinit-time in seconds: 2
Transmit interval in seconds: 30
```

View LLDP global traffic

```
OS10# show lldp traffic
LLDP traffic statistics:
Total Frames Out               : 0
Total Entries Aged             : 0
```

```
Total Frames In                 : 0
Total Frames Received In Error  : 0
Total Frames Discarded          : 0
Total TLVS Unrecognized         : 0
Total TLVs Discarded            : 0
```

View LLDP interface traffic

```
OS10# show lldp traffic interface ethernet 1/1/1
LLDP Traffic Statistics:
Total Frames Out                    : 0
Total Entries Aged                  : 0
Total Frames In                     : 0
Total Frames Received In Error      : 0
Total Frames Discarded              : 0
Total TLVS Unrecognized             : 0
Total TLVs Discarded                : 0

LLDP MED Traffic Statistics:
Total Med Frames Out                : 0
Total Med Frames In                 : 0
Total Med Frames Discarded          : 0
Total Med TLVS Discarded            : 0
Total Med Capability TLVS Discarded: 0
Total Med Policy TLVS Discarded     : 0
Total Med Inventory TLVS Discarded : 0
```

Adjacent agent advertisements

- View brief information about adjacent devices in EXEC mode.
  ```
  show lldp neighbors
  ```
- View all information that neighbors are advertising in EXEC mode.
  ```
  show lldp neighbors detail
  ```
- View all interface-specific information that neighbors are advertising in EXEC mode.
  ```
  show lldp neighbors interface ethernetnode/slot/port[:subport]
  ```

View LLDP neighbors

```
OS10# show lldp neighbors
Loc PortID        Rem Host Name   Rem Port Id         Rem Chassis Id
-------------------------------------------------------------------------
ethernet1/1/2     Not Advertised  fortyGigE 0/56      00:01:e8:8a:fd:35
ethernet1/1/20:1  Not Advertised  GigabitEthernet 1/0 00:01:e8:05:db:05
```

View LLDP neighbors detail

```
OS10# show lldp neighbors interface ethernet 1/1/1 detail

Remote Chassis ID Subtype: Mac address (4)
Remote Chassis ID: 00:13:21:57:ca:40
Remote Port Subtype: Interface name (5)
Remote Port ID: ethernet1/1/10
Remote Port Description: Ethernet port 1
Local Port ID: ethernet1/1/1
Locally assigned remote Neighbor Index: 3
Remote TTL: 120
Information valid for next 105 seconds
Time since last information change of this neighbor: 00:00:15
Remote System Name: LLDP-pkt-gen
Remote Management Address (IPv4): 10.1.1.1
Remote System Desc: LLDP packet generator using scapy
Existing System Capabilities:  Repeater, Bridge, Router
Enabled System Capabilities:  Repeater, Bridge, Router
Remote Max Frame Size: 0
Remote Aggregation Status: false
```

```
MAC PHY Configuration:
    Auto-neg supported: 1
    Auto-neg enabled: 1
    Auto-neg advertised capabilities:
        10BASE-T half duplex mode,
        10BASE-T full duplex mode,
        100BASE-TX half duplex mode,
        100BASE-TX full duplex mode
MED Capabilities:
    Supported:
        LLDP-MED Capabilities,
        Network Policy,
       Location Identification,
        Extended Power via MDI - PSE,
        Extended Power via MDI - PD,
        Inventory Management
    Current:
        LLDP-MED Capabilities,
        Network Policy,
        Location Identification,
        Extended Power via MDI - PD,
        Inventory Management
    Device Class: Endpoint Class 3
Network Policy:
    Application: voice, Tag: Tagged, Vlan: 50, L2 Priority: 6, DSCP Value: 46
Inventory Management:
    H/W Revision  : 12.1.1
    F/W Revision  : 10.1.9750B
    S/W Revision  : 10.1.9750B
    Serial Number : B11G152
    Manufacturer  : Dell
    Model         : S6010-ON
     Asset ID      : E1001
Power-via-MDI:
    Power Type: PD Device
    Power Source: Local and PSE
    Power Priority: Low
    Power required: 6.5
Location Identification:
    Civic-based:
        2C:02:49:4E:01:02:54:4E:03:07:43:68:65:6E:6E:61:69:04:06:47:75:69:
        6E:64:79:05:0B:53:49:44:43:4F:49:6E:64:45:73:74:17:05:4F:54:50:2D:
        31
    ECS-ELIN:
        39:39:36:32:30:33:35:38:32:34
```

View LLDP neighbors interface

```
OS10# show lldp neighbors interface ethernet 1/1/1
Loc PortID          Rem Host Name       Rem Port Id      Rem Chassis Id
-------------------------------------------------------------------------
ethernet1/1/1       OS10                ethernet1/1/2  4:17:eb:f7:06:c4
```

Time to live

The information received from a neighbor expires after a specific amount of time (in seconds) called TTL. The TTL is the LLDPDU transmit interval (hello) and an integer is called a multiplier. For example, LLDPDU transmit interval (30) times the multiplier (4), (30 x 4 = 120). The default multiplier is 4, with a default TTL of 120 seconds.

1 Adjust the TTL value in CONFIGURATION mode.

```
lldp holdtime-multiplier
```

2 Return to the default multiplier value in CONFIGURATION mode.

```
no lldp holdtime-multiplier
```

Configure TTL

```
OS10(config)# lldp holdtime-multiplier 2
```

Return multiplier value

```
OS10(config)# no lldp holdtime-multiplier
```

LLDP commands

clear lldp counters

Clears LLDP and LLDP-MED transmit, receive, and discard statistics from all physical interfaces.

Syntax	`clear lldp counters`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	The counter default value resets to zero for all physical interfaces.
Example	`OS10# clear lldp counters`
Supported Releases	10.2.0E or later

clear lldp table

Clears LLDP neighbor information for all interfaces.

Syntax	`clear lldp table`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	Neighbor information clears on all interfaces.
Example	`OS10# clear lldp table`
Supported Releases	10.2.0E or later

lldp enable

Enables or disables LLDP globally.

Syntax	`lldp enable`
Parameters	None
Default	Enabled
Command Mode	CONFIGURATION

Usage Information	This command enables LLDP globally for all Ethernet PHY interfaces, except on those interfaces where you manually disable LLDP. The `no` version of this command disables LLDP globally irrespective of whether you manually disable LLDP on an interface.
Example	`OS10(config)# lldp enable`
Supported Releases	10.3.1E or later

lldp holdtime-multiplier

Configures the multiplier value for the hold time in seconds.

Syntax	`lldp holdtime-multiplier` *integer*
Parameters	*integer* — Enter the holdtime-multiplier value in seconds, from 2 to 10.
Default	4 seconds
Command Mode	CONFIGURATION
Usage Information	Hold time is the amount of time in seconds that a receiving system waits to hold the information before discarding it. Formula: Hold Time = (Updated Frequency Interval) *X* (Hold Time Multiplier). The `no` version of this command resets the value to the default.
Example	`OS10(config)# lldp holdtime-multiplier 2`
Supported Releases	10.2.0E or later

lldp med fast-start-repeat-count

Configures the number of packets sent during the activation of the fast start mechanism.

Syntax	`lldp-med fast-start-repeat-count` *number*
Parameters	*number* — Enter the number of packets sent during the activation of the fast start mechanism, from 1 to 10.
Default	3
Command Mode	CONFIGURATION
Usage Information	None
Example	`OS10(config)# lldp med fast-start-repeat-count 5`
Supported Releases	10.2.0E or later

lldp med

Enables or disables LLDP-MED on an interface.

Syntax	`lldp med {enable | disable}`
Parameters	• `enable` — Enable LLDP-MED on the interface. • `disable` — Disable LLDP-MED on the interface.
Default	Enabled with network-policy TLV

Command Mode	INTERFACE
Usage Information	LLDP-MED communicates the types of TLVs that the endpoint device and network-connectivity device support. Use the `no lldp med` or `lldp med disable` command to disable LLDP-MED on a specific interface.
Example	`OS10(conf-if-eth1/1/1)# lldp med disable`
Supported Releases	10.2.0E or later

lldp med network-policy

Manually defines an LLDP-MED network policy.

Syntax

```
lldp-med network-policy number app {voice | voice-signaling | guest-voice |
guestvoice-signaling | softphone-voice | streaming-video | video-conferencing |
video-signaling} {vlan vlan-id vlan-type {tag | untag} priority priority dscp
dscp value}
```

Parameters

- *number* — Enter a network policy index number, from 1 to 32.
- app — Enter the type of applications available for the network policy:
  - voice — Voice network-policy application.
  - voice-signaling — Voice-signaling network-policy application.
  - guest-voice — Guest voice network-policy application.
  - guestvoice-signaling — Guest voice signaling network policy application.
  - softphone-voice — SoftPhone voice network-policy application.
  - streaming-video — Streaming video network-policy application.
  - video-conferencing — Voice conference network-policy application.
  - video-signaling — Video signaling network-policy application.
- vlan *vlan-id* — Enter the VLAN number for the selected application, from 1 to 4093.
- vlan-type — Enter the type of VLAN the application uses.
- tag — Enter a tagged VLAN number.
- untag — Enter an untagged VLAN number.
- priority *priority* — Enter the user priority set for the application.
- dscp *dscp value* — Enter the DSCP value set for the application.

Default	Not configured
Command Mode	CONFIGURATION
Usage Information	You can create a maximum of 32 network policies and attach the LLDP-MED network policies to a port.
Example	`OS10(config)# lldp med network-policy 10 app voice vlan 10 vlan-type tag priority 2 dscp 1`
Supported Releases	10.2.0E or later

lldp med network-policy (Interface)

Attaches or removes an LLDP-MED network policy to or from an interface.

Syntax

```
lldp-med network-policy {add | remove} number
```

Parameters

- `add` — Attach the network policy to an interface.
- `remove` — Remove the network policy from an interface.
- *`number`* — Enter a network policy index number, from 1 to 32.

Default Not configured

Command Mode INTERFACE

Usage Information Attach only one network policy for per interface.

Example

```
OS10(conf-if-eth1/1/5)# lldp med network-policy add 1
```

Supported Release 10.2.0E or later

lldp med tlv-select

Configures the LLDP-MED TLV type to transmit or receive.

Syntax

```
lldp med tlv-select {network—policy | inventory}
```

Parameters

- `network-policy` — Enable or disable the port description TLV.
- `inventory` — Enable or disable the system TLV.

Default Enabled

Command Mode INTERFACE

Usage Information None

Example

```
OS10(conf-if-eth1/1/3)# lldp med tlv-select network-policy
```

Supported Releases 10.2.0E or later

lldp port-description-tlv advertise

Specifies whether to advertise the interface description or the port id in the port description TLV.

Syntax

```
lldp port-description-tlv advertise [description | port-id]
```

Parameters

- `description`—Advertise interface description.
- `port-id`—Advertise port id.

Default Interface description is advertised.

Command Mode INTERFACE

Usage Information Determines whether to advertise the interface description or the port id in the port description TLV.

Example

```
OS10(conf-if-eth1/1/1)# lldp port-description-tlv advertise description
OS10(conf-if-eth1/1/1)# lldp port-description-tlv advertise port-id
```

Supported Releases 10.4.3.0 or later

lldp receive

Enables or disables the LLDP packet reception on a specific interface.

Syntax	`lldp receive`
Parameters	None
Default	Not configured
Command Mode	INTERFACE
Usage Information	Enable LLDP globally on the system before using the `lldp receive` command. The `no` version of this command disables the reception of LLDP packets.
Example	`OS10(conf-if-eth1/1/3)# lldp receive`
Supported Releases	10.2.0E or later

lldp reinit

Configures the delay time in seconds for LLDP to initialize on any interface.

Syntax	`lldp reinit` *seconds*
Parameters	*seconds* — Enter the delay timer value in seconds, from 1 to 10.
Default	2 seconds
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(config)# lldp reinit 5`
Supported Releases	10.2.0E or later

lldp timer

Configures the rate in seconds at which LLDP packets send to the peers.

Syntax	`lldp timer` *seconds*
Parameters	*seconds* — Enter the LLDP timer rate in seconds, from 5 to 254.
Default	30 seconds
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command sets the LLDP timer back to its default value.
Example	`OS10(config)# lldp timer 25`
Supported Releases	10.2.0E or later

lldp tlv-select basic-tlv

Enables or disables TLV attributes to transmit and receive LLDP packets.

Syntax

```
lldp tlv-select basic-tlv {port-description | system-name | system-description
| system-capabilities | management-address [ipv4 | ipv6]}
```

Parameters

- `port-description` — Enable or disable the port description TLV.
- `system-name` — Enable or disable the system TLV.
- `system-description` — Enable or disable the system description TLV.
- `system-capabilities` — Enable or disable the system capabilities TLV.
- `management-address` — Enable or disable the management address TLV (IPv4 and IPv6).
- `management-address ipv4` - Enable or disable only the IPv4 management address TLV.
- `management-address ipv6` - Enable or disable only the IPv6 management address TLV.

Default Enabled

Command Mode INTERFACE

Usage Information The `no` form of the command disables TLV attribute transmission and reception in LLDP packets.

Example

```
OS10(conf-if-eth1/1/3)# lldp tlv-select basic-tlv system-name
```

Supported Releases 10.2.0E or later

lldp tlv-select dot1tlv

Enables or disables the dot.1 TLVs to transmit in LLDP packets.

Syntax

```
lldp tlv-select dot1tlv { port-vlan-id | link-aggregation}
```

Parameters

- `port-vlan-id` — Enter the port VLAN ID.
- `link-aggregation` — Enable the link aggregation TLV.

Default Enabled

Command Mode INTERFACE

Usage Information The `lldp tlv-select dot1tlv link-aggregation` command advertises link aggregation as a dot1 TLV in the LLDPDUs. The `no` version of this command disables TLV transmissions.

Example (Port)

```
OS10(conf-if-eth1/1/3)# lldp tlv-select dot1tlv port-vlan-id
```

Example (Link Aggregation)

```
OS10(conf-if-eth1/1/3)# lldp tlv-select dot1tlv link-aggregation
```

Supported Releases 10.2.0E or later

lldp tlv-select dot3tlv

Enables or disables the dot3 TLVs to transmit in LLDP packets.

Syntax `lldp tlv-select dot3tlv {macphy-config | max-framesize}`

Parameters

- `macphy-config` — Enable the port VLAN ID TLV.
- `max-framesize` — Enable maximum frame size TLV.

Default Enabled

Command Mode INTERFACE

Usage Information The `no` version of this command disables TLV transmission.

Example

```
OS10(conf-if-eth1/1/3)# lldp tlv-select dot3tlv macphy-config
```

Supported Releases 10.2.0E or later

lldp transmit

Enables the transmission of LLDP packets on a specific interface.

Syntax `lldp transmit`

Parameters None

Default Not configured

Command Mode INTERFACE

Usage Information The `no` version of this command disables the transmission of LLDP packets on a specific interface.

Example

```
OS10(conf-if-eth1/1/9)# lldp transmit
```

Supported Releases 10.2.0E or later

show lldp interface

Displays the LLDP information advertised from a specific interface.

Syntax `show lldp interface ethernet` *`node/slot/port[:subport]`* `[med | local—device]`

Parameters

- `ethernet` *`node/slot/port[:subport]`* — Enter the Ethernet interface information.
- `med` — Enter the interface to view the MED information.
- `local-device` — Enter the interface to view the local-device information.

Default None

Command Mode EXEC

Usage Information Use the `med` parameter to view MED information for a specific interface. Use the `local-device` parameter to view inventory details.

Example

```
OS10# show lldp interface ethernet 1/1/5
ethernet1/1/5
Tx State             : Enabled
Rx State             : Enabled
Tx SEM State         : initialize
Rx SEM State         : wait-port-operational
Notification Status  : Disabled
Notification Type    : mis-configuration
DestinationMacAddr   : 01:80:c2:00:00:0e
```

Example (Local Device)

```
OS10# show lldp interface ethernet 1/1/1 local-device
Device ID: 00:0c:29:e5:aa:f4
Port ID: ethernet1/1/1
System Name: OS10
Capabilities: Bridge Router
System description: Dell networking Operating system
Port description: Connected to end point device
Time To Live: 120
LLDP MED Capabilities: Capabilities, Network Policy
LLDP MED Device Type: Network connectivity
```

Example (MED)

```
OS10# show lldp interface ethernet 1/1/20:1 med
Port             |Capabilities|Network Policy|Location|Inventory|POE
-----------------|------------|--------------|--------|---------|---
ethernet1/1/20:1|         Yes|           Yes|      No|       No| No
Network Polices :
```

Supported Releases 10.2.0E or later

show lldp errors

Displays the LLDP errors related to memory allocation failures, queue overflows, and table overflows.

Syntax `show lldp errors`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show lldp errors
Total Memory Allocation Failures: 0
Total Input Queue Overflows: 0
Total Table Overflows: 0
```

Supported Release 10.2.0E or later

show lldp med

Displays the LLDP MED information for all the interfaces.

Syntax `show lldp med`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information Use the `show lldp interface` command to view MED information for a specific interface.

Example

```
OS10# show lldp med
Fast Start Repeat Count: 3
LLDP MED Device Type: Network Connectivity
Port            |Capabilities|Network Policy|Location|Inventory|POE
----------------|------------|--------------|--------|---------|---
ethernet1/1/1   |         Yes|           Yes|      No|       No| No
ethernet1/1/2   |         Yes|           Yes|      No|       No| No
ethernet1/1/3   |         Yes|           Yes|      No|       No| No
ethernet1/1/4   |         Yes|           Yes|      No|      Yes| No
ethernet1/1/5   |         Yes|           Yes|      No|       No| No
ethernet1/1/6   |         Yes|           Yes|      No|       No| No
ethernet1/1/7   |         Yes|           Yes|      No|      Yes| No
ethernet1/1/8   |         Yes|           Yes|      No|       No| No
ethernet1/1/9   |         Yes|           Yes|      No|       No| No
ethernet1/1/10  |         Yes|           Yes|      No|       No| No
ethernet1/1/11  |         Yes|           Yes|      No|       No| No
ethernet1/1/12  |         Yes|           Yes|      No|       No| No
ethernet1/1/13  |         Yes|           Yes|      No|       No| No
ethernet1/1/14  |         Yes|           Yes|      No|       No| No
ethernet1/1/15  |         Yes|           Yes|      No|       No| No
ethernet1/1/16  |         Yes|           Yes|      No|       No| No
ethernet1/1/17  |         Yes|           Yes|      No|       No| No
ethernet1/1/18  |         Yes|           Yes|      No|       No| No
ethernet1/1/19  |         Yes|           Yes|      No|       No| No
ethernet1/1/20  |         Yes|           Yes|      No|       No| No
ethernet1/1/21  |         Yes|           Yes|      No|       No| No
ethernet1/1/22  |         Yes|           Yes|      No|       No| No
ethernet1/1/23  |         Yes|           Yes|      No|       No| No
ethernet1/1/24  |         Yes|           Yes|      No|       No| No
ethernet1/1/25  |         Yes|           Yes|      No|       No| No
ethernet1/1/26  |         Yes|           Yes|      No|       No| No
ethernet1/1/27  |         Yes|           Yes|      No|       No| No
ethernet1/1/28  |         Yes|           Yes|      No|       No| No
ethernet1/1/29  |         Yes|           Yes|      No|       No| No
ethernet1/1/30  |         Yes|           Yes|      No|       No| No
ethernet1/1/31  |         Yes|           Yes|      No|       No| No
ethernet1/1/32  |         Yes|           Yes|      No|       No| No
```

Supported Releases 10.2.0E or later

show lldp neighbors

Displays the status of the LLDP neighbor system information.

Syntax `show lldp neighbors [detail | interface ethernet node/slot/port[:subport]]`

Parameters

- `detail` — View LLDP neighbor detailed information.
- `interface ethernet node/slot/port[:subport]` — Enter the Ethernet interface information.

Command Mode EXEC

Usage Information This command status information includes local port ID, remote host name, remote port ID, and remote node ID.

Example

```
OS10# show lldp neighbors
Loc PortID        Rem Host Name   Rem Port Id         Rem Chassis Id
-------------------------------------------------------------------
ethernet1/1/2     Not Advertised  fortyGigE 0/56      00:01:e8:8a:fd:35
ethernet1/1/20:1  Not Advertised  GigabitEthernet 1/0 00:01:e8:05:db:05
```

Example (Detail)

```
OS10# show lldp neighbors interface ethernet 1/1/1 detail
```

```
Remote Chassis ID Subtype: Mac address (4)
Remote Chassis ID: 00:50:56:a6:29:54
Remote Port Subtype: Interface alias (1)
Remote Port ID: ethernet1/1/1
Remote Port Description: ethernet1/1/1
Local Port ID: ethernet1/1/1
Locally assigned remote Neighbor Index: 2
Remote TTL: 120
Information valid for next 99 seconds
Time since last information change of this neighbor: 15:51:41
Remote System Name: OS10
Remote System Desc: OS10
Existing System Capabilities:  Repeater, Bridge, Router
Enabled System Capabilities:  Repeater, Bridge, Router
Remote Port Vlan ID: 1
Remote Max Frame Size: 1532
Remote Aggregation Status: false
MAC PHY Configuration:
    Auto-neg supported: 1
    Auto-neg enabled: 1
    Auto-neg advertised capabilities:
        1000BASE-T half duplex mode
Dell EMC Organization Specific Detail:
Originator: Switch
    Service Tag: B8D1XC2
    Product Base: base1
    Product Serial Number: sn1
    Product Part Number: pn1
```

Example (Interface)

```
OS10# show lldp neighbors interface ethernet 1/1/1
Loc PortID          Rem Host Name       Rem Port Id      Rem Chassis Id
-------------------------------------------------------------------------
ethernet1/1/1       OS10                ethernet1/1/2  4:17:eb:f7:06:c4
```

Supported Releases 10.2.0E or later

show lldp timers

Displays the LLDP hold time, delay time, and update frequency interval configuration information.

Syntax `show lldp timers`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show lldp timers
LLDP Timers:
Holdtime in seconds: 120
Reinit-time in seconds: 6
Transmit interval in seconds: 30
```

Supported Releases 10.2.0E or later

show lldp tlv-select interface

Displays the TLVs enabled for an interface.

Syntax `show lldp tlv-select interface ethernet node/slot/port[:subport]`

Parameters `ethernet node/slot/port[:subport]` — Enter the Ethernet interface information, from 1 to 253.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show lldp tlv-select interface ethernet 1/1/4
port-description
system-name
system-description
system-cabalities
management-address
port-vlan
mac-phy-config
link-aggregation
max-frame-size
```

Supported Releases 10.2.0E or later

show lldp traffic

Displays LLDP traffic information including counters, packets transmitted and received, discarded packets, and unrecognized TLVs.

Syntax `show lldp traffic [interface ethernet node/slot/port[:subport]]`

Parameters `interface ethernet node/slot/port[:subport]` — (Optional) Enter the Ethernet interface information to view the LLDP traffic.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show lldp traffic
LLDP Traffic Statistics:
Total Frames Out                  : 1504
Total Entries Aged                : 2
Total Frames In                   : 67
Total Frames Received In Error    : 0
Total Frames Discarded            : 0
Total TLVS Unrecognized           : 0
Total TLVs Discarded              : 0
```

Example (Interface)

```
OS10# show lldp traffic interface ethernet 1/1/2
LLDP Traffic Statistics:
Total Frames Out                  : 45
Total Entries Aged                : 1
Total Frames In                   : 33
Total Frames Received In Error    : 0
Total Frames Discarded            : 0
Total TLVS Unrecognized           : 0
Total TLVs Discarded              : 0
```

```
LLDP MED Traffic Statistics:
Total Med Frames Out               : 2
Total Med Frames In                : 1
Total Med Frames Discarded         : 0
Total Med TLVS Discarded           : 0
Total Med Capability TLVS Discarded: 0
Total Med Policy TLVS Discarded    : 0
Total Med Inventory TLVS Discarded : 0
```

Supported Releases 10.2.0E or later

show nework-policy profile

Displays the network policy profiles.

Syntax `show network-policy profile [profile number]`

Parameters *profile number* — (Optional) Enter the network policy profile number, from 1 to 32.

Default Not configured

Command Mode EXEC

Usage Information If you do not enter the network profile ID, all configured network policy profiles display.

Example

```
OS10# show network-policy profile 10
Network Policy Profile 10
   voice vlan 17 cos 4
  Interface:
   none
Network Policy Profile 30
   voice vlan 30 cos 5
  Interface:
   none
Network Policy Profile 36
   voice vlan 4 cos 3
  Interface:
       ethernet 1/1/1,ethernet 1/1/3-5
```

Supported Releases 10.2.0E or later

Media Access Control

All Ethernet switching ports maintain media access control (MAC) address tables. Each physical device in your network contains a MAC address. OS10 devices automatically enter learned MAC addresses as dynamic entries in the MAC address table.

Learned MAC address entries are subject to aging. Set the aging timer to zero (0) to disable MAC aging. For any dynamic entry, if no packet arrives on the device with the MAC address as the source or destination address within the timer period, the address is removed from the table.

- Enter an aging time (in seconds) in CONFIGURATION mode, from 0 to 1000000, default 1800.
  `mac address-table aging-time seconds`

Configure Aging Time

```
OS10(config)# mac address-table aging-time 900
```

Disable Aging Time

```
OS10(config)# mac address-table aging-time 0
```

Static MAC Address

You manually configure a static MAC address entry. A static entry is not subject to aging.

- Create a static MAC address entry in the MAC address table in CONFIGURATION mode.

```
mac-address-table static nn:nn:nn:nn:nn vlan vlan-id interface [ethernet node/slot/
port[:subport] | port-channel channel-number]
```

Set Static MAC Address

```
OS10(config)# mac address-table static 34:17:eb:f2:ab:c6 vlan 10 interface ethernet 1/1/5
```

MAC Address Table

OS10 maintains a list of MAC address table entries.

- View the contents of the MAC address table in EXEC mode.

```
show mac address-table {dynamic | static} [address mac-address | vlan vlan-id | interface
{ethernet node/slot/port[:subport] | port-channel number}] [count [vlan vlan-id] [interface
{type node/slot/port[:subport] | port-channel number}]
```

  - `dynamic` — (Optional) Displays dynamic MAC address table entry information.
  - `static` — (Optional) Displays static MAC address table entry information.
  - `address mac-address` — (Optional) Displays MAC address information.
  - `interface ethernet node/slot/port[:subport]` — (Optional) Displays a list of dynamic and static MAC address entries.
  - `interface port-channel number` — (Optional) Displays port channel information, from 1 to 128.
  - `count` — (Optional) Displays the number of dynamic and static MAC address entries.
  - `vlan vlan-id` — (Optional) Displays information for a specified VLAN only, from 1 to 4093.

View MAC Address Table Entries

```
OS10# show mac address-table
VlanId  Mac Address             Type            Interface
1       00:00:15:c6:ca:49       dynamic         ethernet1/1/21
1       00:00:20:2a:25:55       dynamic         ethernet1/1/21
1       90:b1:1c:f4:aa:ce       dynamic         ethernet1/1/21
1       90:b1:1c:f4:aa:c6       dynamic         ethernet1/1/21
10      34:17:eb:02:8c:33       static          ethernet1/1/1
```

View MAC Address Table Count

```
OS10# show mac address-table count
MAC Entries for all vlans :
Dynamic Address Count :                 4
Static Address (User-defined) Count :   1
Total MAC Addresses in Use:             5
```

Clear MAC Address Table

You can clear dynamic address entries that in the MAC address table maintains.

- Clear the MAC address table of dynamic entries in EXEC mode.

```
clear mac address-table dynamic [[all] [address mac_addr] [vlan vlan-id] [interface {ethernet
type node/slot/port[:subport] | port-channel number}]
```

- `all` — (Optional) Clear all dynamic entries.
- `address mac_address` — (Optional) Clear a MAC address entry.
- `vlan vlan-id` — (Optional) Clear a MAC address table entry from a VLAN number, from 1 to 4093.
- `ethernet node/slot/port[:subport]` — (Optional) Clear an Ethernet interface entry.
- `port–channel number` — (Optional) Clear a port-channel number, from 1 to 128.

Clear MAC Address Table

```
OS10# clear mac address-table dynamic vlan 20 interface ethernet 1/2/20
```

MAC Commands

clear mac address-table dynamic

Clears L2 dynamic address entries from the MAC address table.

Syntax

```
clear mac address-table dynamic {all | address mac_addr | vlan vlan-id |
interface {ethernet node/slot/port[:subport] | port-channel number}}
```

Parameters

- `all` — (Optional) Delete all MAC address table entries.
- `address mac_addr` — (Optional) Delete a configured MAC address from the address table in nn:nn:nn:nn:nn:nn format.
- `vlan vlan-id` — (Optional) Delete all entries based on the VLAN number from the address table, from 1 to 4093.
- `interface` — (Optional) Clear the interface type:
  - `ethernet node/slot/port[:subport]` — Delete the Ethernet interface configuration from the address table.
  - `port-channel channel-number` — Delete the port-channel interface configuration from the address table, from 1 to 128.

Default Not configured

Command Mode EXEC

Usage Information Use the `all` parameter to remove all dynamic entries from the address table.

Example

```
OS10# clear mac address-table dynamic all
```

Example (VLAN)

```
OS10# clear mac address-table dynamic vlan 20
```

Supported Releases 10.2.0E or later

mac address-table aging-time

Configures the aging time for entries in the L2 address table.

Syntax `mac address-table aging-time seconds`

Parameters `seconds` — Enter the aging time for MAC table entries in seconds, from 0 to 1000000.

Default 1800 seconds

Command Mode	CONFIGURATION
Usage Information	Set the aging timer to zero (`0`) to disable MAC address aging for all dynamic entries. The aging time counts from the last time that the device detected the MAC address.
Example	`OS10(config)# mac address-table aging-time 3600`
Supported Releases	10.2.0E or later

mac address-table static

Configures a static entry for the L2 MAC address table.

Syntax `mac address-table static mac-address vlan vlan-id interface {ethernet node/slot/port[:subport] | port-channel number}`

Parameters

- `mac-address` — Enter the MAC address to add to the table in nn:nn:nn:nn:nn:nn format.
- `vlan vlan-id` — Enter the VLAN to apply the static MAC address to, from 1 to 4093.
- `interface` — Enter the interface type:
  - `ethernet node/slot/port[:subport]` — Enter the Ethernet information.
  - `port-channel channel-number` — Enter a port-channel interface number, from 1 to 128.

Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command resets the value to the default.
Example (VLAN)	`OS10(config)# mac address-table static 34:17:eb:f2:ab:c6 vlan 1 interface ethernet 1/1/30`
Example (Port-Channel)	`OS10(config)# mac address-table static 34:17:eb:02:8c:33 vlan 10 interface port-channel 1`
Supported Releases	10.2.0E or later

show mac address-table

Displays information about the MAC address table.

Syntax `show mac address-table [address mac-address | aging-time | [count [vlan vlan-id] | dynamic | interface {ethernet node/slot/port[:subport] | port-channel number}]| static [address mac-address] | vlan vlan-id`

Parameters

- `address mac-address` — (Optional) Displays MAC address table information.
- `aging-time` — (Optional) Displays MAC address table aging-time information.
- `count` — (Optional) Displays the number of dynamic and static MAC address entries.
- `dynamic` — (Optional) Displays dynamic MAC address table entries only.
- `interface` — Set the interface type:
  - `ethernet node/slot/port[:subport]` — Displays MAC address table information for a physical interface.

  - `port-channel` *`channel-number`* — Displays MAC address table information for a port-channel interface, from 1 to 128.
- `static` — (Optional) Displays static MAC address table entries only.
- `vlan` *`vlan-id`* — (Optional) Displays VLAN information only, from 1 to 4093.

Default Not configured

Command Mode EXEC

Usage Information The network device maintains static MAC address entries saved in the startup configuration file, and reboots and deletes dynamic entries.

Example (Address)

```
OS10# show mac address-table address 90:b1:1c:f4:a6:8f
VlanId  Mac Address             Type           Interface
1       90:b1:1c:f4:a6:8f       dynamic         ethernet1/1/3
```

Example (Aging Time)

```
OS10# show mac address-table aging-time
Global Mac-address-table aging time : 1800
```

Example (Count)

```
OS10# show mac address-table count
MAC Entries for all vlans :
Dynamic Address Count : 5
Static Address (User-defined) Count : 0
Total MAC Addresses in Use: 5
```

Example (Dynamic)

```
OS10# show mac address-table dynamic
VlanId  Mac Address          Type     Interface
1       90:b1:1c:f4:a6:8f    dynamic  ethernet1/1/3
```

Example (Ethernet)

```
OS10# show mac address-table interface ethernet 1/1/3
VlanId  Mac Address         Type     Interface
1       66:38:3a:62:31:3a   dynamic  ethernet1/1/3
```

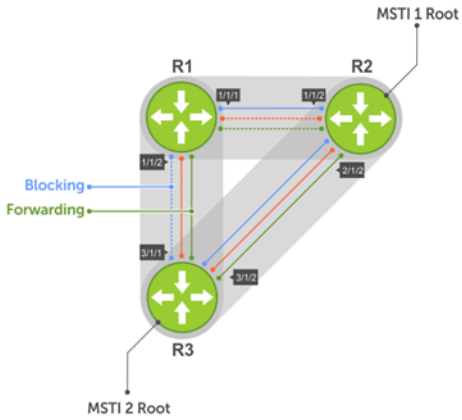
Supported Releases 10.2.0E or later

Multiple Spanning-Tree

MST is a rapid spanning-tree protocol based on a spanning-tree variation that improves on per-VLAN RPVST+. You can configure Multiple Spanning-Tree Instances (MSTIs) and map multiple VLANs to one spanning-tree instance to reduce the total number of instances. RPVST+ allows a spanning-tree instance for each VLAN. This 1:1 approach is not suitable if you have multiple VLANs — each spanning-tree instance costs bandwidth and processing resources.

When you enable MST, all ports in Layer 2 (L2) mode participate in MST. OS10 only supports one MST region.

Achieve load balancing using the MST protocol (MSTP). When you map three VLANs to two multiple spanning tree instances (MSTIs), VLAN 100 traffic takes a different path than VLAN 200 and 300 traffic.



Configuring MST is a four-step process:

1 Enable MST, if the current running spanning tree protocol (STP) version is not MST.
2 (Optional) Map the VLANs to different instances to achieve load balancing.
3 Ensure the same region name is configured in all the bridges running MST.
4 (Optional) Configure the revision number.

Configure MSTP

When you enable MST globally, all L2 physical, port-channel, and VLAN interfaces automatically assign to MSTI zero (0). Within an MSTI, only one path from one bridge to another is enabled for forwarding.

- Enable MST in CONFIGURATION mode.
  ```
  spanning-tree mode mst
  ```

Configure and verify MSTP

```
OS10(config)# spanning-tree mode mst
OS10(config)# do show spanning-tree
show spanning-tree mst configuration
Region Name: ravi
Revision: 0
MSTI    VID
0       1,7-4093
1       2
2       3
3       4
4       5
5       6
```

Add or remove interfaces

By default, all interfaces are enabled in L2 switchport mode, and all L2 interfaces are part of spanning-tree.

- Disable spanning-tree on an interface in INTERFACE mode.
  ```
  spanning-tree disable
  ```
- Enable MST on an interface in INTERFACE mode.
  ```
  no spanning-tree disable
  ```

Create instances

You can create multiple MSTP instances and map VLANs. A single MSTI provides no more benefit than RSTP. To take full advantage of the MSTP, create multiple MSTIs and map VLANs to them.

1 Enter an instance number in CONFIGURATION mode.

```
spanning tree mst configuration
```

2 Enter the MST instance number in MULTIPLE-SPANNING-TREE mode, from 0 to 63.

```
instance instance-number
```

3 Enter the VLAN and IDs to participate in the MST instance in MULTIPLE-SPANNING-TREE mode, from 1 to 4096.

```
instance vlan-id
```

Create MST instances

```
OS10(config)# spanning-tree mst configuration
OS10(conf-mst)# name force10
OS10(conf-mst)# revision 100
OS10(conf-mst)# instance 1 vlan 2-10
OS10(conf-mst)# instance 2 vlan 11-20
OS10(conf-mst)# instance 3 vlan 21-30
```

View VLAN instance mapping

```
OS10# show spanning-tree mst configuration
Region Name: force10
Revision: 100
MSTI    VID
0       1,31-4093
1       2-10
2       11-20
3       21-30
```

View port forwarding/discarding state

```
OS10# show spanning-tree msti 0 brief
Spanning tree enabled protocol msti with force-version mst
MSTI 0 VLANs mapped 1,31-4093
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32768, Address 3417.4455.667f
Root Bridge hello time 2, max age 20, forward delay 15, max hops 20
Bridge ID    Priority 32768, Address 90b1.1cf4.a523
Configured hello time 2, max age 20, forward delay 15, max hops 20
CIST regional root ID Priority 32768, Address 90b1.1cf4.a523
CIST external path cost 500
Interface                                                 Designated
Name             PortID  Prio  Cost      Sts   Cost  Bridge ID    PortID
-------------------------------------------------------------------------------
ethernet1/1/1   128.260  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.260
ethernet1/1/2   128.264  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.264
ethernet1/1/3   128.268  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.268
ethernet1/1/4   128.272  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.272
ethernet1/1/5   128.276  128   500       FWD   0    32768   3417.4455.667f  128.146
ethernet1/1/6   128.280  128   500       BLK   0    32768   3417.4455.667f  128.150
ethernet1/1/7   128.284  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.284
ethernet1/1/8   128.288  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.288
ethernet1/1/9   128.292  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.292
ethernet1/1/10  128.296  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.296
ethernet1/1/11  128.300  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.300
ethernet1/1/12  128.304  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.304
ethernet1/1/13  128.308  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.308
ethernet1/1/14  128.312  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.312
ethernet1/1/15  128.316  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.316
ethernet1/1/16  128.320  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.320
```

```
ethernet1/1/17  128.324  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.324
ethernet1/1/18  128.328  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.328
ethernet1/1/19  128.332  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.332
ethernet1/1/20  128.336  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.336
ethernet1/1/21  128.340  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.340
ethernet1/1/22  128.344  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.344
ethernet1/1/23  128.348  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.348
ethernet1/1/24  128.352  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.352
ethernet1/1/25  128.356  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.356
ethernet1/1/26  128.360  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.360
ethernet1/1/27  128.364  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.364
ethernet1/1/28  128.368  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.368
ethernet1/1/29  128.372  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.372
ethernet1/1/30  128.376  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.376
ethernet1/1/31  128.380  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.380
ethernet1/1/32  128.384  128   200000000 BLK   0    32768   90b1.1cf4.a523  128.384
Interface
Name              Role   PortID    Prio  Cost       Sts   Cost  Link-type Edge
-------------------------------------------------------------------------------
ethernet1/1/1     Disb   128.260   128   200000000 BLK   0     AUTO      No
ethernet1/1/2     Disb   128.264   128   200000000 BLK   0     AUTO      No
ethernet1/1/3     Disb   128.268   128   200000000 BLK   0     AUTO      No
ethernet1/1/4     Disb   128.272   128   200000000 BLK   0     AUTO      No
ethernet1/1/5     Root   128.276   128   500       FWD   0     AUTO      No
ethernet1/1/6     Altr   128.280   128   500       BLK   0     AUTO      No
ethernet1/1/7     Disb   128.284   128   200000000 BLK   0     AUTO      No
ethernet1/1/8     Disb   128.288   128   200000000 BLK   0     AUTO      No
ethernet1/1/9     Disb   128.292   128   200000000 BLK   0     AUTO      No
ethernet1/1/10    Disb   128.296   128   200000000 BLK
```

Root selection

MSTP determines the root bridge according to the lowest bridge ID. Assign a lower bridge priority to increase its likelihood of becoming the root bridge.

- Assign a bridge priority number to a specific instance in CONFIGURATION mode, from 0 to 61440 in increments of 4096, default 32768. Use a lower priority number to increase the likelihood of the bridge to become a root bridge.

  ```
  spanning-tree mst instance-number priority priority
  ```

Assign root bridge priority

```
OS10(config)# spanning-tree mst 0
```

Verify root bridge priority

```
OS10# show spanning-tree active
Spanning tree enabled protocol msti with force-version mst
MSTI 0 VLANs mapped 1,31-4093
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32768, Address 3417.4455.667f
Root Bridge hello time 2, max age 20, forward delay 15, max hops 20
Bridge ID    Priority 32768, Address 90b1.1cf4.a523
Configured hello time 2, max age 20, forward delay 15, max hops 20
CIST regional root ID Priority 32768, Address 90b1.1cf4.a523
CIST external path cost 500
Interface                                          Designated
Name           PortID   Prio  Cost  Sts  Cost Bridge ID     PortID
-------------------------------------------------------------------------------
ethernet1/1/5  128.276  128   500   FWD  0    32768   3417.4455.667f  128.146
ethernet1/1/6  128.280  128   500   BLK  0    32768   3417.4455.667f  128.150
Interface
Name           Role PortID   Prio Cost Sts  Cost Link-type Edge
-------------------------------------------------------------
ethernet1/1/5  Root 128.276  128  500  FWD  0    AUTO      No
ethernet1/1/6  Altr 128.280  128  500  BLK  0    AUTO      No
```

Non-Dell EMC hardware

OS10 supports only one MST region. For a bridge to be in the same MST region as another, the three unique name, revision, and VLAN-to-instance-mapping attributes must match. The default values for the name and revision number match on all Dell EMC hardware. If you have non-Dell EMC hardware that participates in MST, ensure these values match on all devices.

A region is a combination of three unique attributes:

- Name — A mnemonic string you assign to the region. The default is the system MAC address.
- Revision — A 2-byte number. The default is 0.
- VLAN-to-instance mapping — Placement of a VLAN in an MSTI.

Region name or revision

You can change the MSTP region name or revision.

- Change the region name in MULTIPLE-SPANNING-TREE mode. A maximum of 32 characters.
  ```
  name name
  ```
- Change the region revision number in MULTIPLE-SPANNING-TREE mode, from 0 to 65535, default 0.
  ```
  revision number
  ```

Configure and verify region name

```
OS10(conf-mstp)# name my-mstp-region
OS10(conf-mstp)# do show spanning-tree mst config
MST region name: my-mstp-region
Revision: 0
MSTI   VID
  1    100
  2    200-300
```

Modify parameters

The root bridge sets the values for forward-delay, hello-time, max-age, and max-hops and overwrites the values set on other MST bridges.

Forward-time	Time an interface waits in the Discarding state and Learning state before it transitions to the Forwarding state.
Hello-time	Interval in which the bridge sends MST BPDUs.
Max-age	Length of time the bridge maintains configuration information before it refreshes that information by recomputing the MST topology.
Max-hops	Maximum number of hops a BPDU travels before a receiving device discards it.

NOTE: Dell EMC recommends that only experienced network administrators change MST parameters. Poorly planned modification of MST parameters can negatively affect network performance.

1 Change the forward-time parameter in CONFIGURATION mode, from 4 to 30, default 15.
```
spanning-tree mst forward-time seconds
```

2 Change the hello-time parameter in CONFIGURATION mode, from 1 to 10, default 2. Dell EMC recommends increasing the hello-time for large configurations, especially configurations with more ports.
```
spanning-tree mst hello-time seconds
```

3 Change the max-age parameter in CONFIGURATION mode, from 6 to 40, default 20.
```
spanning-tree mst max-age seconds
```

4 Change the max-hops parameter in CONFIGURATION mode, from 1 to 40, default 20.

```
spanning-tree mst max-hops number
```

MST configuration

```
OS10(config)# spanning-tree mst
OS10(config)# spanning-tree mst forward-time 16
OS10(config)# spanning-tree mst hello-time 5
OS10(config)# spanning-tree mst max-age 10
OS10(config)# spanning-tree mst max-hops 30
```

View MSTP parameter values

```
OS10# show spanning-tree active
Spanning tree enabled protocol msti with force-version mst
MSTI 0 VLANs mapped 1,31-4093
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32768, Address 3417.4455.667f
Root Bridge hello time 2, max age 20, forward delay 15, max hops 20
Bridge ID    Priority 32768, Address 90b1.1cf4.a523
Configured hello time 10, max age 40, forward delay 30, max hops 40
CIST regional root ID Priority 32768, Address 90b1.1cf4.a523
CIST external path cost 500
Interface                                                  Designated
Name           PortID   Prio  Cost Sts   Cost Bridge ID          PortID
-------------------------------------------------------------------------------
ethernet1/1/5  128.276  128   500  FWD   0    32768    3417.4455.667f  128.146
ethernet1/1/6  128.280  128   500  BLK   0    32768    3417.4455.667f  128.150
Interface
Name            Role  PortID   Prio  Cost  Sts  Cost  Link-type Edge
-----------------------------------------------------------------
ethernet1/1/5   Root  128.276  128   500   FWD  0     AUTO      No
ethernet1/1/6   Altr  128.280  128   500   BLK  0     AUTO      No
```

Interface parameters

Adjust two interface parameters to increase or decrease the likelihood that a port becomes a forwarding port.

Port cost	Interface type value. The greater the port cost, the less likely the port is a forwarding port.
Port priority	Influences the likelihood that a port is selected as a forwarding port if several ports have the same port cost.

Default values for the port cost by interface:

- 100-Mb/s Ethernet interfaces — 200000
- 1-Gigabit Ethernet interfaces — 20000
- 10-Gigabit Ethernet interfaces — 2000
- Port-channel with 100 Mb/s Ethernet interfaces — 180000
- Port-channel with 1-Gigabit Ethernet interfaces — 18000
- Port-channel with 10-Gigabit Ethernet interfaces — 1800

1 Change the port cost of an interface in INTERFACE mode, from 0 to 200000000.

```
spanning-tree msti number cost cost
```

2 Change the port priority of an interface in INTERFACE mode, from 0 to 240 in increments of 16, default 128.

```
spanning-tree msti number priority priority
```

View MSTi interface configuration

```
OS10(conf-if-eth1/1/7)# do show spanning-tree msti 0 interface ethernet 1/1/7
ethernet1/1/7 of MSTI 0 is Designated Forwarding
Edge port: No (default)
Link type: point-to-point (auto)
```

```
Boundary: Yes, Bpdu-filter: Disable, Bpdu-Guard: Disable, Shutdown-on-Bpdu-Guard-violation: No
Root-Guard: Disable, Loop-Guard: Disable
Bpdus (MRecords) Sent: 69, Received: 0
Interface                                    Designated
Name           PortID  Prio  Cost  Sts  Cost Bridge ID         PortID
-------------------------------------------------------------------------
ethernet1/1/7  0.284   0     1     FWD  0    32768  90b1.1cf4.9b8a 0.284
```

EdgePort Forward traffic

EdgePort allows the interface to forward traffic approximately 30 seconds sooner as it skips the Blocking and Learning states. The `spanning-tree bpduguard enable` command causes the interface hardware to shut down when it receives a BPDU.

CAUTION: Configure EdgePort only on links connecting to an end station. EdgePort can cause loops if you enable it on an interface connected to a network.

When you implement BPDU guard, although the interface is placed in the Error Disabled state when receiving the BPDU, the physical interface remains in the Up state. The hardware discards regular network traffic after a BPDU violation. BPDUs forward to the CPU, where they are discarded as well.

- Enable EdgePort on an interface in INTERFACE mode.

  ```
  spanning-tree port type edge
  ```

Configure EdgePort

```
OS10(conf-if-eth1/1/4)# spanning-tree port type edge
```

View interface status

```
OS10# show spanning-tree interface ethernet 1/1/4
ethernet1/1/4 of MSTI 0 is designated Forwarding
Edge port:yes port guard :none (default)
Link type is point-to-point (auto)
Boundary: YES bpdu filter :disable bpdu guard :disable  bpduguard shutdown-on-
violation :disable  RootGuard:  disable LoopGuard  disable
Bpdus (MRecords) sent 610, received 5
Interface                                          Designated
Name           PortID   Prio  Cost Sts  Cost  Bridge ID              PortID
-------------------------------------------------------------------------
ethernet1/1/4  128.272  128   500  FWD  0     32768    90b1.1cf4.a911 128.272
=========================================================================
```

Spanning-tree extensions

STP extensions provide a means to ensure efficient network convergence by securely enforcing the active network topology. OS10 supports BPDU filtering, BPDU guard, root guard, and loop guard STP extensions.

BPDU filtering	Protects the network from unexpected flooding of BPDUs from an erroneous device. Enabling BPDU Filtering instructs the hardware to drop BPDUs and prevents flooding from reaching the CPU. BPDU filtering is enabled by default on Edge ports. All BPDUs received on the Edge port drops. If you explicitly configure BPDU filtering on a port, that port drops all BPDUs it receives.
BPDU guard	Blocks the L2 bridged ports and LAG ports connected to end hosts and servers from receiving any BPDUs. When you enable BPDU guard, it places a bridge or LAG port in the Error_Disable or Blocking state if the port receives any BPDU frames. In a LAG, all member ports, including new members are placed in the Blocking state. The network traffic drops but the port continues to forward BPDUs to the CPU that are later dropped. To prevent further reception of BPDUs, configure a port to shut down using the `shutdown` command. The port can only resume operation from the Shutdown state after manual intervention.

Root guard	Avoids bridging loops and preserves the root bridge position during network transitions. STP selects the root bridge with the lowest priority value. During network transitions, another bridge with a lower priority may attempt to become the root bridge and cause unpredictable network behavior. To avoid such an attempt and preserve the position of the root bridge, configure the `spanning-tree guard root` command. Root guard is enabled on ports that are designated ports. The root guard configuration applies to all VLANs configured on the port.
Loop guard	Prevents L2 forwarding loops caused by a cable or interface hardware failure. When a hardware failure occurs, a participating spanning tree link becomes unidirectional and a port stops receiving BPDUs. When a blocked port stops receiving BPDUs, it transitions to a Forwarding state causing spanning tree loops in the network. Enable loop guard on a port that transitions to the Loop-Inconsistent state until it receives BPDUs using the `spanning-tree guard loop` command. After BPDUs are received, the port moves out of the Loop-Inconsistent or blocking state and transitions to an appropriate state determined by STP. Enabling loop guard on a per-port basis enables it on all VLANs configured on the port. If you disable loop guard on a port, it moves to the Listening state.

If you enable BPDU filter and BPDU guard on the same port, the BPDU filter configuration takes precedence. Root guard and Loop guard are mutually exclusive. Configuring one overwrites the other from the active configuration.

1 Enable spanning-tree BPDU filter in INTERFACE mode.

```
spanning-tree bpdufilter enable
```

- To shut down the port channel interface, all member ports are disabled in the hardware.
- To add a physical port to a port-channel already in the Error Disable state, the new member port is also disabled in the hardware.
- To remove a physical port from a port-channel in Error Disable state, the Error Disabled state clears on this physical port. The physical port is enabled in the hardware.

To clear the Error Disabled state:

- Use the `shutdown` command on the interface.
- Use the `spanning-tree bpdufilter disable` command to disable the BPDU guard on the interface.
- Use the `spanning-tree disable` command to disable STP on the interface.

2 Enable STP BPDU guard in INTERFACE mode.

```
spanning-tree bpduguard enable
```

- To shut down the port channel interface, all member ports are disabled in the hardware.
- To add a physical port to a port-channel already in the Error Disable state, the new member port is also disabled in the hardware.
- To remove a physical port from a port-channel in Error Disable state, the Error Disabled state clears on this physical port. The physical port is enabled in the hardware.

To clear the Error Disabled state:

- Use the `shutdown` command on the interface.
- Use the `spanning-tree bpduguard disable` command to disable the BPDU guard on the interface.
- Use the `spanning-tree disable` command to disable STP on the interface.

3 Set the guard types to avoid loops in INTERFACE mode.

```
spanning-tree guard {loop | root | none}
```

- `loop` — Set the guard type to loop.
- `none` — Set the guard type to none.
- `root` — Set the guard type to root.

BPDU filter

```
OS10(conf-if-eth1/1/4)# spanning-tree bpdufilter enable
OS10(conf-if-eth1/1/4)# do show spanning-tree interface ethernet 1/1/4
ethernet1/1/4 of vlan1 is designated Blocking
Edge port:no (default) port guard :none (default)
Link type is point-to-point (auto)
Boundary: NO  bpdu filter : Enable bpdu guard :  bpduguard shutdown-on-
violation :disable  RootGuard:  enable LoopGuard  disable
Bpdus (MRecords) sent 134, received 138
Interface                                          Designated
```

```
Name           PortID   Prio Cost Sts  Cost  Bridge ID          PortID
-------------------------------------------------------------------------
ethernet1/1/4  128.272  128  500  BLK  500   32769    90b1.1cf4.a911 128.272
```

BPDU guard

```
OS10(config)# interface ethernet 1/1/4
OS10(conf-if-eth1/1/4)# spanning-tree bpduguard enable
OS10(conf-if-eth1/1/4)# do show spanning-tree interface ethernet 1/1/4
ethernet1/1/4 of vlan1 is designated Blocking
Edge port:no (default) port guard :none (default)
Link type is point-to-point (auto)
Boundary: NO  bpdu filter : Enable bpdu guard :  bpduguard shutdown-on-
violation :enable  RootGuard:  enable LoopGuard  disable
Bpdus (MRecords) sent 134, received 138
Interface                                                Designated
Name           PortID  Prio Cost Sts  Cost Bridge ID         PortID
-----------------------------------------------------------------
ethernet1/1/4  128.272 128  500  BLK  500  32769    90b1.1cf4.a911 128.272
```

Loop guard

```
OS10(config)# interface ethernet 1/1/4
OS10(conf-if-eth1/1/4)# spanning-tree guard loop
OS10(conf-if-eth1/1/4)# do show spanning-tree interface ethernet 1/1/4
ethernet1/1/4 of vlan1 is root Forwarding
Edge port:no (default) port guard :none (default)
Link type is point-to-point (auto)
Boundary: NO  bpdu filter : bpdu guard :  bpduguard shutdown-on-
violation :disable  RootGuard:  disable LoopGuard  enable
Bpdus (MRecords) sent 7, received 20
Interface                                                 Designated
Name           PortID  Prio  Cost Sts  Cost Bridge ID           PortID
-------------------------------------------------------------------------
ethernet1/1/4  128.272 128   500  FWD  0    32769    90b1.1cf4.9d3b 128.272
```

Root guard

```
OS10(conf-if-eth1/1/4)# spanning-tree guard root
OS10(conf-if-eth1/1/4)# do show spanning-tree interface ethernet 1/1/4
ethernet1/1/4 of vlan1 is root Forwarding
Edge port:no (default) port guard :none (default)
Link type is point-to-point (auto)
Boundary: NO  bpdu filter : bpdu guard :  bpduguard shutdown-on-
violation :disable  RootGuard:  enable LoopGuard  disable
Bpdus (MRecords) sent 7, received 33
Interface                                                 Designated
Name           PortID  Prio Cost Sts Cost Bridge ID        PortID
-----------------------------------------------------------------
ethernet1/1/4  128.272 128  500  BLK 500  32769  90b1.1cf4.a911 128.272
```

Recover BPDU guard error disabled ports

When OS10 detects a BPDU guard violation for an STP enabled port, the system shuts the port down. Use the BPDU Guard error disable recovery option to recover the shut-down ports automatically.

1 When there is BPDU guard violation on a port, OS10 either shuts down the port or moves it to BLOCKED state. Use the following command in CONFIGURATION mode, that shuts down the port. The `no` version of the command moves the port to BLOCKED state.

```
errdisable detect cause bpduguard
```

2 To delay the recovery of ports when there is a BPDU Guard violation, configure the recovery interval timer in CONFIGURATION mode, .

```
errdisable recovery interval interval-value
```

3 In CONFIGURATION mode, use the following command to enable recovering of ports shut down due to detection of a BPDU Guard violation. When the recovery option is enabled, the port is brought up after the recovery timer expires. When the recovery option is

disabled, the port remains shut down indefinitely. You must manually bring up the port using the `shutdown` and `no shutdown` commands. The `no` version of the command disables the recovery option.

```
errdisable recovery cause bpduguard
```

Example configuration

```
OS10(config)# errdisable detect cause bpduguard
OS10(config)# errdisable recovery interval 45
OS10(config)# errdisable recovery cause bpduguard
```

View detect and recovery details

```
OS10# show errdisable detect

Error-Disable Cause              Detect Status
-------------------------------------------
bpduguard                        Enabled
```

```
OS10# show errdisable recovery

Error-Disable Recovery Timer Interval: 300 seconds

Error-Disable Reason             Recovery Status
-------------------------------------------------
bpduguard                        Enabled

                                                        Recovery Time left
Interface              Errdisable Cause                 (seconds)
-----------------------------------------------------------------------
ethernet 1/1/1:1       bpduguard                         273
ethernet 1/1/2         bpduguard                         4
port-channel 12        bpduguard                         45
```

Setting spanning-tree link type for rapid state transitions

To fasten the spanning-tree state transitions, you can set the link type to point-to-point. To set the link type to point-to-point:

- Use the following command in INTERFACE mode.

  ```
  spanning-tree link-type point-to-point
  ```

MST commands

errdisable detect cause bpduguard

Configures the port to be shut down or moves the port to blocked state on detecting a BPDU guard violation.

Syntax	`errdisable detect cause bpduguard`
Parameters	None
Default	Enabled
Command Mode	CONFIGURATION
Usage Information	This command applies only to STP-enabled ports. The command takes effect only when BPDU guard is configured on a port. When the detect cause option is enabled, the port is shut down whenever there is a BPDU guard violation.

When the option is disabled, the port is not shut down but moved to BLOCKING state whenever there is a BPDU guard violation. In this case, the port is operationally DOWN in spanning-tree mode and when the recovery timer expires, the port is UP irrespective of the recovery cause configuration.

The `no` version of the command disables the detect cause option.

Example

```
OS10(config)# errdisable detect cause bpduguard
```

Supported Releases 10.4.2.0 or later

errdisable recovery cause bpduguard

Enables to recover the ports shut down due to BPDU Guard violation.

Syntax `errdisable recovery cause bpduguard`

Parameters None

Default Disabled

Command Mode CONFIGURATION

Usage Information This command applies only to STP-enabled ports. The command takes effect only when BPDU guard is configured on a port and `errdisable detect cause bpduguard`is enabled on the port.

When the recovery option is enabled, the port is brought up after the recovery timer expires.

When the recovery option is disabled, the port is shut down indefinitely. You must manually bring up the port using the `shutdown` and `no shutdown` commands.

The `no` version of the command disables the recovery option.

Example

```
OS10(config)# errdisable recovery cause bpduguard
```

Supported Releases 10.4.2.0 or later

errdisable recovery interval

Configures recovery interval timer to delay the recovery of ports when there is a BPDU Guard violation.

Syntax `errdisable recovery interval` *`interval-value`*

Parameters *`interval-value`*—Enter the time interval in seconds. The range is from 30 to 65535.

Default 300 seconds

Command Mode CONFIGURATION

Usage Information This command applies only to STP-enabled ports. The command takes effect only when BPDU guard is configured on a port.

The recovery timer starts whenever there is a BPDU guard violation.

The `no` version of the command resets the timer to default value.

Example

```
OS10(config)# errdisable recovery interval 45
```

Supported Releases 10.4.2.0 or later

instance

Configures MST instances and one or multiple VLANs mapped to the MST instance.

Syntax `instance instance-number {vlan vlan-range}`

Parameters

- *instance* — Enter an MST instance value, from 0 to 63.
- `vlan` *range* — Enter a VLAN range value, from 1 to 4093.

Default Not configured

Command Mode MULTIPLE-SPANNING-TREE

Usage Information By default, all VLANs map to MST instance zero (0) unless you are using the `vlan` *range* command to map the VLANs to a non-zero instance. The `no` version of this command removes the instance-related configuration.

Example

```
OS10(conf-mst)# instance 1 vlan 2-10
OS10(conf-mst)# instance 2 vlan 11-20
OS10(conf-mst)# instance 3 vlan 21-30
```

Supported Releases 10.2.0E or later

name

Assigns a name to the MST region.

Syntax `name region-name`

Parameters `region-name` — Enter a name for an MST region. A maximum of 32 characters.

Default System MAC address

Command Mode MULTIPLE-SPANNING-TREE

Usage Information By default, the MST protocol assigns the system MAC address as the region name. Two MST devices within the same region must share the same region name, including matching case.

Example

```
OS10(conf-mst)# name my-mst-region
```

Supported Releases 10.2.0E or later

revision

Configures a revision number for the MSTP configuration.

Syntax `revision number`

Parameters *number* — Enter a revision number for the MSTP configuration, from 0 to 65535.

Default 0

Command Mode MULTIPLE-SPANNING-TREE

Usage Information	To have a bridge in the same MST region as another, the default values for the revision number must match on all Dell EMC hardware devices. If there are non-Dell EMC devices, ensure the revision number value matches on all the devices. For more information, see Non-Dell Hardware.
Example	`OS10(conf-mst)# revision 10`
Supported Releases	10.2.0E or later

spanning-tree bpdufilter

Enables or disables BPDU filtering on an interface.

Syntax	`spanning-tree bpdufilter {enable | disable}`
Parameters	• `enable` — Enables the BPDU filter on an interface. • `disable` — Disables the BPDU filter on an interface.
Default	Disabled
Command Mode	INTERFACE
Usage Information	Use the `enable` parameter to enable BPDU filtering.
Example	`OS10(conf-if-eth1/1/4)# spanning-tree bpdufilter enable`
Supported Releases	10.2.0E or later

spanning-tree bpduguard

Enables or disables the BPDU guard on an interface.

Syntax	`spanning-tree bpduguard {enable | disable}`
Parameters	• `enable` — Enables the BPDU guard filter on an interface. • `disable` — Disables the BPDU guard filter on an interface.
Default	Disabled
Command Mode	INTERFACE
Usage Information	BPDU guard prevents a port from receiving BPDUs. If the port receives a BPDU, it is placed in the Error-Disabled state.
Example	`OS10(conf-if-eth1/1/4)# spanning-tree bpduguard enable`
Supported Releases	10.2.0E or later

spanning-tree disable

Disables Spanning-Tree mode configured with the `spanning-tree mode` command globally on the switch or on specified interfaces.

Syntax	`spanning-tree disable`
Parameters	None

Default	Not configured.
Usage Information	The `no` version of this command re-enables STP and applies the currently configured spanning-tree settings.
Command Mode	CONFIGURATION INTERFACE
Example	`OS10(config)# interface ethernet 1/1/4` `OS10(config-if-eth1/1/4)# spanning-tree disable`
Supported Releases	10.3.0E or later

spanning-tree guard

Enables or disables loop guard or root guard on an interface.

Syntax	`spanning-tree guard {loop | root | none}`
Parameters	• `loop` — Enables loop guard on an interface. • `root` — Enables root guard on an interface. • `none` — Sets the guard mode to none.
Default	Not configured
Usage Information	Root guard and loop guard configurations are mutually exclusive. Configuring one overwrites the other from the active configuration.
Command Mode	INTERFACE
Example	`OS10(conf-if-eth1/1/4)# spanning-tree guard root`
Supported Releases	10.2.0E or later

spanning-tree link-type

Sets the spanning-tree link type for faster convergence.

Syntax	`spanning-tree link-type {auto | point-to-point | shared}`
Parameters	• `auto` — Enter the keyword to sets the link type based on the duplex setting of the interface. • `point-to-point`—Specifies that the interface is a point-to-point or full-duplex link. • `shared`—Specifies that the interface is a half-duplex medium.
Default	Auto
Command Mode	INTERFACE
Usage Information	To enable fast state transition, use the `point-to-point` option.
Example	`OS10(config)# spanning-tree link-type point-to-point`
Supported Releases	10.4.3.0 or later

spanning-tree mode

Enables an STP type: RSTP, Rapid-PVST+, or MST.

Syntax `spanning-tree mode {rstp | mst | rapid-pvst}`

Parameters

- `rstp` — Sets STP mode to RSTP.
- `mst` — Sets STP mode to MST.
- `rapid-pvst` — Sets STP mode to RPVST+.

Default RPVST+

Command Mode CONFIGURATION

Usage Information All STP instances stop in the previous STP mode, and restart in the new mode. You can also change to RSTP/MST mode.

Example (RSTP)

```
OS10(config)# spanning-tree mode rstp
```

Example (MST)

```
OS10(config)# spanning-tree mode mst
```

Supported Releases 10.2.0E or later

spanning-tree mst

Configures an MST instance and determines root and bridge priorities.

Syntax `spanning-tree mst` *`instance number`* `priority | root {primary | secondary}`

Parameters

- *`instance number`* — Enter an MST instance number, from 0 to 63.
- `priority` *`priority value`* — Set a bridge priority value in increments of 4096, from 0 to 61440. Valid priority values are: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. All other values are rejected.
- `root` — Enter a primary or secondary root.
- `primary` — Enter a device as a primary root.
- `secondary` — Enter a device as a secondary root.

Default Not configured

Command Mode CONFIGURATION

Usage Information The MSTP determines the root bridge but you can assign one bridge a lower priority to increase the probability it being the root bridge. A lower *`priority-value`* increases the probability of the bridge becoming a root bridge. The `no` version of this command resets the value to the default.

Example

```
OS10(config)# spanning-tree mst 0 priority 0
OS10(config)# spanning-tree mst 2 root primary
```

Supported Releases 10.2.0E or later

spanning-tree msti

Configures the MSTI, cost, and priority values for an interface.

Syntax `spanning-tree msti instance {cost cost | priority value}`

Parameters

- `msti instance` — Enter the MST instance number, from 0 to 63.
- `cost cost` — (Optional) Enter a port cost value, from 1 to 200000000. Default values:
  - 100 Mb/s Ethernet interface = 200000
  - 1-Gigabit Ethernet interface = 20000
  - 10-Gigabit Ethernet interface = 2000
  - Port-channel interface with one 100 Mb/s Ethernet = 200000
  - Port-channel interface with one 1 Gigabit Ethernet = 20000
  - Port-channel interface with one 10 Gigabit Ethernet = 2000
  - Port-channel with two 1 Gigabit Ethernet = 18000
  - Port-channel with two 10 Gigabit Ethernet = 1800
  - Port-channel with two 100 Mbps Ethernet = 180000
- `priority value` — Enter a value in increments of 16 as the priority, from 0 to 240, default 128.

Default Priority value is 128

Command Mode INTERFACE

Usage Information The `cost` value is based on the interface type. The greater the `cost` value, the less likely the port is selected to be a forwarding port. The `priority` influences the likelihood that a port is selected to be a forwarding port if several ports have the same cost value.

Example

```
OS10(conf-if-eth1/1/1)# spanning-tree msti 1 priority 0
OS10(conf-if-eth1/1/1)# spanning-tree msti 1 cost 3
```

Supported Releases 10.2.0E or later

spanning-tree mst configuration

Enters MST mode to configure MSTP from Configuration mode.

Syntax `spanning-tree mst configuration`

Parameters None

Default Disabled

Command Mode CONFIGURATION

Usage Information Use this command to enter STP MST configuration mode.

Example

```
OS10(config)# spanning-tree mst configuration
OS10(conf-mst)#
```

Supported Releases 10.2.0E or later

spanning-tree mst disable

Disables spanning tree on the specified MST instance.

Syntax	`spanning-tree mst` *instance-number* `disable`
Parameters	*instance-number*—Enter the instance number, from 0 to 63.
Default	Enabled
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command enables spanning tree on the specified MST instance.
Example	`OS10(config)# spanning-tree mst 10 disable`
Supported Releases	10.4.0E(R1) or later

spanning-tree mst force-version

Configures a forced version of STP to transmit BPDUs.

Syntax	`spanning-tree mst force-version {stp | rstp}`
Parameters	• `stp` — Forces the version for the BPDUs transmitted by MST to STP. • `rstp` — Forces the version for the BPDUs transmitted by MST to RSTP.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	Forces a bridge that supports MST to operate in a STP-compatible mode.
Example	`OS10(config)# spanning-tree mst force-version`
Supported Releases	10.2.0E or later

spanning-tree mst forward-time

Configures a time interval for the interface to wait in the Blocking state or the Learning state before moving to the Forwarding state.

Syntax	`spanning-tree mst forward-time` *seconds*
Parameters	*seconds* — Enter the number of seconds an interface waits in the Blocking or Learning States before moving to the Forwarding state, from 4 to 30.
Default	15 seconds
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(config)# spanning-tree mst forward-time 16`
Supported Releases	10.2.0E or later

spanning-tree mst hello-time

Sets the time interval between generation and transmission of MSTP BPDUs.

Syntax	`spanning-tree mst hello-time` *seconds*
Parameters	*seconds* — Enter a hello-time interval value in seconds, from 1 to 10.
Default	2 seconds
Command Mode	CONFIGURATION
Usage Information	Dell EMC recommends increasing the hello-time for large configurations, especially configurations with multiple ports. The `no` version of this command resets the value to the default.
Example	`OS10(config)# spanning-tree mst hello-time 5`
Supported Releases	10.2.0E or later

spanning-tree mst mac-flush-threshold

Configures the mac-flush threshold value for a specific instance.

Syntax	`spanning-tree mst` *instance-number* `mac-flush-threshold` *threshold-value*
Parameters	• *instance-number*—Enter the instance number, from 0 to 4094. • *threshold-value*—Enter the threshold value for the number of flushes, from 0 to 65535. The default value is 5.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	This threshold indicates the number of port-based flush requests allowed to be invoked before starting the flush optimization. When the flush interval value is non-zero, port-and-instance-based flushing is triggered until the threshold is reached. Once the threshold is reached, instance-based flushing is triggered and the timer starts. The `no` form of the command sets the flush indication threshold of the specific instance to its default value.
Example	`OS10(config)# spanning-tree mst 10 mac-flush-threshold 255`
Supported Releases	10.4.0E(R1) or later

spanning-tree mst max-age

Configures the time period the bridge maintains configuration information before refreshing the information by recomputing the MST topology.

Syntax	`max-age` *seconds*
Parameters	*seconds* — Enter a maximum age value in seconds, from 6 to 40.
Default	20 seconds
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command resets the value to the default.

Example	`OS10(config)# spanning-tree mst max-age 10`
Supported Releases	10.2.0E or later

spanning-tree mst max-hops

Configures the maximum hop count for a BPDU to travel before it is discarded.

Syntax	`spanning-tree mst max-hops` *number*
Parameters	*number* — Enter a maximum hop value, from 6 to 40.
Default	20
Command Mode	CONFIGURATION
Usage Information	A device receiving BPDUs waits until the `max-hops` value expires before discarding it. When a device receives the BPDUs, it decrements the received value of the remaining hops and uses the resulting value as remaining-hops in the BPDUs. If the remaining MSTP 1333 hops reach zero, the device discards the BPDU and ages out any information that it holds for the port. The command configuration applies to all common IST (CIST) in the MST region.
Example	`OS10(config)# spanning-tree mst max-hops 30`
Supported Releases	10.2.0E or later

spanning-tree port

Sets the port type as the EdgePort.

Syntax	`spanning-tree port type edge`
Parameters	None
Default	Not configured
Command Mode	INTERFACE
Usage Information	When you configure an EdgePort on a device running STP, the port immediately transitions to the Forwarding state. Only configured ports connected to end hosts act as EdgePorts.
Example	`OS10(config)# spanning-tree port type edge`
Supported Releases	10.2.0E or later

show errdisable

Displays information on errdisable configurations and port recovery status.

Syntax	`show errdisable [detect | recovery]`
Parameters	• `detect`—Displays details of detect cause configuration. • `recovery`—Displays details of recovery cause, recovery interval, and recovery status of the error disabled port.
Default	None

Command Mode EXEC

Usage Information None

Example

```
OS10# show errdisable detect

Error-Disable Cause               Detect Status
-------------------------------------------
bpduguard                          Enabled
```

```
OS10# show errdisable recovery

Error-Disable Recovery Timer Interval: 300 seconds

Error-Disable Reason               Recovery Status
-------------------------------------------------
bpduguard                          Enabled

                                                        Recovery Time left
Interface               Errdisable Cause                (seconds)
-----------------------------------------------------------------------
ethernet 1/1/1:1        bpduguard                        273
ethernet 1/1/2          bpduguard                        4
port-channel 12         bpduguard                        45
```

Supported Releases 10.4.2.0 or later

show spanning-tree mst

Displays MST configuration information.

Syntax `show spanning-tree mst configuration`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information Enable MSTI before using this command.

Example

```
OS10# show spanning-tree mst configuration
Region Name: asia
Revision: 0
MSTI    VID
0       1,7-4093
1       2
2       3
3       4
4       5
5       6
```

Supported Releases 10.2.0E or later

show spanning-tree msti

Displays MST instance information.

Syntax `show spanning-tree msti [instance-number [brief | guard | virtual-interface | interface interface]]`

Parameters

- *instance-number* — (Optional) Displays MST instance information, from 0 to 63.

- `brief` — (Optional) Displays MST instance summary information.
- `guard` — (Optional) Displays which guard is enabled and the current port state.
- `virtual-interface`—(Optional) Displays MST information specific to VLT.
- `interface` *`interface`*—(Optional) Displays interface type information:
  - `ethernet` *`node/slot/port[:subport]`* — Enter the Ethernet port information, from 1 to 48.
  - `port-channel` — Enter the port-channel interface information, from 1 to 128.

Default Not configured

Command Mode EXEC

Usage Information View the MST instance information for a specific MST instance number in detail or brief, or view physical Ethernet ports or port-channel information.

Example (Brief)

```
OS10# show spanning-tree msti 0 brief
Spanning tree enabled protocol msti with force-version mst
MSTI 0 VLANs mapped 1-99,101-199,301-4093
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32768, Address 90b1.1cf4.9b8a
Root Bridge hello time 2, max age 20, forward delay 15, max hops 20
Bridge ID    Priority 32768, Address 90b1.1cf4.9b8a
We are the root of MSTI 0
Configured hello time 2, max age 20, forward delay 15, max hops 20
Interface                                              Designated
Name          PortID   Prio  Cost    Sts  Cost Bridge ID   PortID
-----------------------------------------------------------------
ethernet1/1/1  132.128 128 200000000 BLK  0 32768 90b1.1cf4.9b8a 128.132
ethernet1/1/2  136.128 128 200000000 BLK  0 32768 90b1.1cf4.9b8a 128.136
ethernet1/1/3 140.128  128 200000000 BLK  0 32768 90b1.1cf4.9b8a 128.140
ethernet1/1/4  144.128 128 200000000 BLK  0 32768 90b1.1cf4.9b8a 128.144
ethernet1/1/5  148.128 128 200000000 BLK  0 32768 90b1.1cf4.9b8a 128.148
ethernet1/1/6  152.128 128 200000000 BLK  0 32768 90b1.1cf4.9b8a 128.152
ethernet1/1/7  156.128 128 200000000 BLK  0 32768 90b1.1cf4.9b8a 128.156
...
Interface
Name           Role  PortID  Prio  Cost     Sts  Cost Link-type Edge
-------------------------------------------------------------------
ethernet1/1/1  Disb 128.132  128  200000000 BLK  0    SHARED    No
ethernet1/1/2  Disb 128.136  128  200000000 BLK  0    SHARED    No
ethernet1/1/3  Disb 128.140  128  200000000 BLK  0    SHARED    No
ethernet1/1/4  Disb 128.144  128  200000000 BLK  0    SHARED    No
ethernet1/1/5  Disb 128.148  128  200000000 BLK  0    SHARED    No
ethernet1/1/6  Disb 128.152  128  200000000 BLK  0    SHARED    No
ethernet1/1/7  Disb 128.156  128  200000000 BLK  0    SHARED    No
ethernet1/1/8  Disb 128.160  128  200000000 BLK  0    SHARED    No
ethernet1/1/9  Disb 128.164  128  200000000 BLK  0    SHARED    No
```

Example (Interface)

```
OS10# show spanning-tree msti 1 interface ethernet 1/1/1
ethernet1/1/1 of vlan1 is root Forwarding
Edge port:no (default) port guard :none (default)
Link type is point-to-point (auto)
Boundary :internal bpdu filter : bpdu guard :  bpduguard shutdown-on-
violation :disable  RootGuard:  disable LoopGuard  disable
Bpdus (MRecords) sent 3779, received 7
Interface                                          Designated
 Name       PortID  Prio  Cost  Sts Cost Bridge ID     PortID
-----------------------------------------------------------------
ethernet1/1/1 128.132 128 20000 FWD 0 32768 74e6.e2f5.dd80 128.132
```

Example (Guard)

```
OS10# show spanning-tree msti 1 guard
Interface
Name             Instance  Sts   Guard Type
-----------------------------------------
ethernet1/1/1    MSTI 1      FWD   root
```

```
ethernet1/1/2     MSTI 1       FWD    loop
ethernet1/1/3     MSTI 1       BLK    none
ethernet1/1/4     MSTI 1       FWD    none
ethernet1/1/5     MSTI 1       BLK    none
ethernet1/1/6     MSTI 1       BLK    none
ethernet1/1/7     MSTI 1       BLK    none
ethernet1/1/8     MSTI 1       BLK    none
...
```

Example (virtual-interface)

```
agg-6146 # show spanning-tree msti 0 virtual-interface
VFP(VirtualFabricPort) of MSTI 0 is Designated Forwarding
Edge port: No (default)
Link type: point-to-point (auto)
Boundary: No, Bpdu-filter: Disable, Bpdu-Guard: Disable, Shutdown-on-Bpdu-Guard-violation: No
Root-Guard: Disable, Loop-Guard: Disable
Bpdus (MRecords) Sent: 250, Received: 240
Interface                                                                                    Designated
Name                          PortID    Prio      Cost      Sts         Cost      Bridge ID
PortID
-------------------------------------------------------------------------------------------------------------------
VFP(VirtualFabricPort)        0.1       0         1         FWD         0         32768     3417.ebf2.a8c4  0.1
```

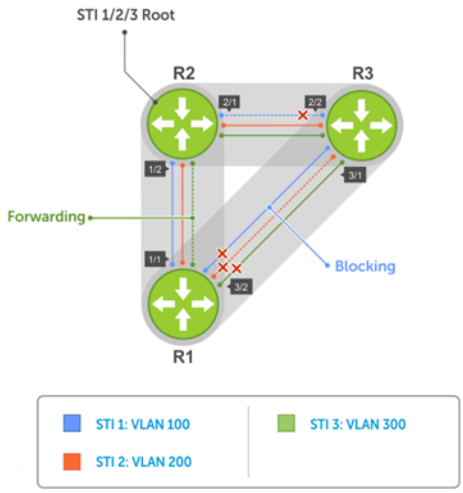
Command History 10.2.0E or later

Rapid per-VLAN spanning-tree plus

RPVST+ is an RSTP to create a single topology per VLAN. RPVST+ is enabled by default, provides faster convergence, and runs on the default VLAN (VLAN 1).

Configuring Rapid-PVST+ is a four-step process:

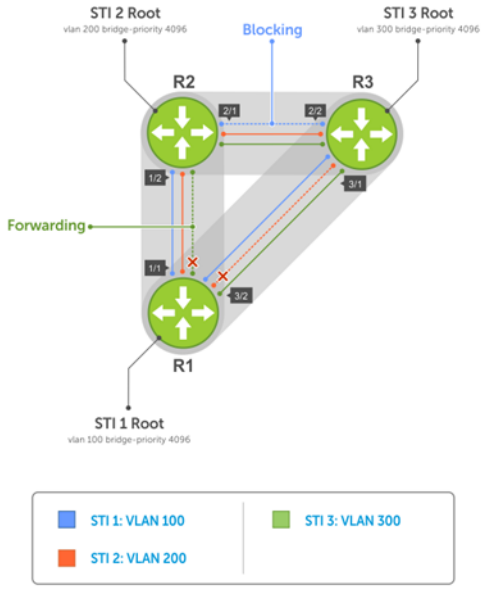
1. Ensure the interfaces are in L2 mode.
2. Place the interfaces in VLANs. By default, switchport interfaces are members of the default (VLAN1).
3. Enable Rapid-PVST+. This step is only required if another variation of STP is present.
4. (Optional) Select a non-default bridge-priority for the VLAN for load balancing.



By default, each VLAN instance is assigned default bridge priority 32768. For example, all three instances have the same forwarding topology. Traffic load balancing is not achievable with this kind of priority assignment. To achieve load balancing, you must assign each instance a different priority, as shown in *Load Balancing with RPVST+*.

Load balance and root selection

All VLANs use the same forwarding topology — R2 is elected as the root and all 10G Ethernet ports have the same cost. RPVST+ changes the bridge priority of each bridge so that a different forwarding topology generates for each VLAN.



To achieve RPVST+ load balancing, assign a different priority on each bridge.

Enable RPVST+

By default, RPVST+ is enabled and creates an instance only after you add the first member port to a VLAN. To participate in RPVST+, port-channel or physical interfaces must be a member of a VLAN. Add all physical and port-channel interfaces to the default VLAN (VLAN1).

- Enable Rapid-PVST+ mode in CONFIGURATION mode.

```
spanning-tree mode rapid-pvst
```

Configure RPVST+

```
OS10(config)# spanning-tree mode rapid-pvst
```

View RPVST+ configuration

```
OS10# show spanning-tree active
Spanning tree enabled protocol rapid-pvst with force-version rstp
VLAN 1
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32768, Address 3417.4455.667f
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 32769, Address 90b1.1cf4.a523
Configured hello time 2, max age 20, forward delay 15
Interface                                              Designated
Name           PortID   Prio Cost Sts  Cost Bridge ID     PortID
-------------------------------------------------------------------------
ethernet1/1/5  128.276  128  500  FWD  0    32768  3417.4455.667f  128.146
ethernet1/1/6  128.280  128  500  BLK  0    32768  3417.4455.667f  128.150
Interface
Name           Role  PortID   Prio Cost Sts  Cost Link-type Edge
```

```
-------------------------------------------------------------
ethernet1/1/5  Root  128.276  128  500  FWD  0    AUTO      No
ethernet1/1/6  Altr  128.280  128  500  BLK  0    AUTO      No
```

Select root bridge

RPVST+ determines the root bridge. Assign one bridge a lower priority to increase the likelihood that it becomes the root bridge. The `show spanning-tree brief` command displays information about all ports regardless of the operational status.

- Assign a number as the bridge priority or designate it as the root in CONFIGURATION mode, from 0 to 61440.

  `spanning-tree {vlan vlan-id priority priority-value}`

  - *vlan-id* — Enter a value between 1 to 4093.
  - priority *priority-value* — Enter the priority value in increments of 4096, default is 32768. The lower the number assigned, the more likely this bridge becomes the root bridge. The bridge priority the valid values are: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, or 61440. All other values are rejected.

Configure root bridge

```
OS10(config)# spanning-tree vlan 1 priority 4096
```

View active configuration

```
OS10(config)# do show spanning-tree active
Spanning tree enabled protocol rapid-pvst with force-version rstp
VLAN 1
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 4097, Address 90b1.1cf4.a523
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 4097, Address 90b1.1cf4.a523
We are the root of VLAN 1
Configured hello time 2, max age 20, forward delay 15
Interface                                              Designated
Name            PortID   Prio Cost Sts  Cost Bridge ID     PortID
-------------------------------------------------------------------
ethernet1/1/5   128.276  128  500  FWD  0    4097   90b1.1cf4.a523  128.276
ethernet1/1/6   128.280  128  500  FWD  0    4097   90b1.1cf4.a523  128.280
Interface
Name            Role  PortID   Prio Cost Sts  Cost Link-type Edge
-------------------------------------------------------------
ethernet1/1/5   Desg  128.276  128  500  FWD  0    AUTO      No
ethernet1/1/6   Desg  128.280  128  500  FWD  0    AUTO      No
```

View brief configuration

```
OS10# show spanning-tree brief
Spanning tree enabled protocol rapid-pvst with force-version rstp
VLAN 1
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 4097, Address 90b1.1cf4.a523
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 4097, Address 90b1.1cf4.a523
We are the root of VLAN 1
Configured hello time 2, max age 20, forward delay 15
Interface                                                    Designated
Name             PortID   Prio  Cost      Sts   Cost  Bridge ID             PortID
-----------------------------------------------------------------------------------
ethernet1/1/1    128.260  128   200000000 FWD   0     32769  0000.0000.0000  128.260
ethernet1/1/2    128.264  128   200000000 FWD   0     32769  0000.0000.0000  128.264
ethernet1/1/3    128.268  128   200000000 FWD   0     32769  0000.0000.0000  128.268
ethernet1/1/4    128.272  128   200000000 FWD   0     32769  0000.0000.0000  128.272
ethernet1/1/5    128.276  128   500       FWD   0     4097   90b1.1cf4.a523  128.276
ethernet1/1/6    128.280  128   500       FWD   0     4097   90b1.1cf4.a523  128.280
ethernet1/1/7    128.284  128   200000000 FWD   0     32769  0000.0000.0000  128.284
ethernet1/1/8    128.288  128   200000000 FWD   0     32769  0000.0000.0000  128.288
ethernet1/1/9    128.292  128   200000000 FWD   0     32769  0000.0000.0000  128.292
```

```
ethernet1/1/10   128.296  128   200000000 FWD   0     32769  0000.0000.0000  128.296
ethernet1/1/11   128.300  128   200000000 FWD   0     32769  0000.0000.0000  128.300
ethernet1/1/12   128.304  128   200000000 FWD   0     32769  0000.0000.0000  128.304
ethernet1/1/13   128.308  128   200000000 FWD   0     32769  0000.0000.0000  128.308
ethernet1/1/14   128.312  128   200000000 FWD   0     32769  0000.0000.0000  128.312
ethernet1/1/15   128.316  128   200000000 FWD   0     32769  0000.0000.0000  128.316
ethernet1/1/16   128.320  128   200000000 FWD   0     32769  0000.0000.0000  128.320
ethernet1/1/17   128.324  128   200000000 FWD   0     32769  0000.0000.0000  128.324
ethernet1/1/18   128.328  128   200000000 FWD   0     32769  0000.0000.0000  128.328
ethernet1/1/19   128.332  128   200000000 FWD   0     32769  0000.0000.0000  128.332
ethernet1/1/20   128.336  128   200000000 FWD   0     32769  0000.0000.0000  128.336
ethernet1/1/21   128.340  128   200000000 FWD   0     32769  0000.0000.0000  128.340
ethernet1/1/22   128.344  128   200000000 FWD   0     32769  0000.0000.0000  128.344
ethernet1/1/23   128.348  128   200000000 FWD   0     32769  0000.0000.0000  128.348
ethernet1/1/24   128.352  128   200000000 FWD   0     32769  0000.0000.0000  128.352
ethernet1/1/25   128.356  128   200000000 FWD   0     32769  0000.0000.0000  128.356
ethernet1/1/26   128.360  128   200000000 FWD   0     32769  0000.0000.0000  128.360
ethernet1/1/27   128.364  128   200000000 FWD   0     32769  0000.0000.0000  128.364
ethernet1/1/28   128.368  128   200000000 FWD   0     32769  0000.0000.0000  128.368
ethernet1/1/29   128.372  128   200000000 FWD   0     32769  0000.0000.0000  128.372
ethernet1/1/30   128.376  128   200000000 FWD   0     32769  0000.0000.0000  128.376
ethernet1/1/31   128.380  128   200000000 FWD   0     32769  0000.0000.0000  128.380
ethernet1/1/32   128.384  128   200000000 FWD   0     32769  0000.0000.0000  128.384
Interface
Name             Role   PortID   Prio  Cost       Sts   Cost Link-type Edge
-------------------------------------------------------------------------
ethernet1/1/1    Disb   128.260  128   200000000 FWD   0    AUTO      No
ethernet1/1/2    Disb   128.264  128   200000000 FWD   0    AUTO      No
ethernet1/1/3    Disb   128.268  128   200000000 FWD   0    AUTO      No
ethernet1/1/4    Disb   128.272  128   200000000 FWD   0    AUTO      No
ethernet1/1/5    Desg   128.276  128   500       FWD   0    AUTO      No
ethernet1/1/6    Desg   128.280  128   500       FWD   0    AUTO      No
ethernet1/1/7    Disb   128.284  128   200000000 FWD   0    AUTO      No
ethernet1/1/8    Disb   128.288  128   200000000 FWD   0    AUTO      No
ethernet1/1/9    Disb   128.292  128   200000000 FWD   0    AUTO      No
ethernet1/1/10   Disb   128.296  128   200000000 FWD   0    AUTO      No
ethernet1/1/11   Disb   128.300  128   200000000 FWD   0    AUTO      No
```

Root assignment

RPVST+ assigns the root bridge according to the lowest bridge ID. Assign one bridge as the root bridge and the other as a secondary root bridge.

- Configure the device as the root or secondary root in CONFIGURATION mode.

  ```
  spanning-tree vlan vlan-id root {primary | secondary}
  ```

  - *vlan-id* — Enter the VLAN ID number, from 1 to 4093.
  - primary — Enter the bridge as primary or root bridge. The primary bridge value is 24576.
  - secondary — Enter the bridge as secondary or secondary root bridge. The secondary bridge value is 28672.

Configure root bridge as primary

```
OS10(config)# spanning-tree vlan 1 root primary
```

Verify root bridge information

```
OS10# show spanning-tree active

Spanning tree enabled protocol rapid-pvst with force-version rstp
VLAN 1
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 24577, Address 90b1.1cf4.a523
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 24577, Address 90b1.1cf4.a523
We are the root of VLAN 1
```

```
Configured hello time 2, max age 20, forward delay 15
Interface                                                  Designated
Name            PortID   Prio  Cost Sts  Cost Bridge ID           PortID
-------------------------------------------------------------------------
ethernet1/1/5   128.276  128   500  FWD  0    24577  90b1.1cf4.a523  128.276
ethernet1/1/6   128.280  128   500  LRN  0    24577  90b1.1cf4.a523  128.280
Interface
Name            Role PortID  Prio Cost Sts  Cost Link-type Edge
-----------------------------------------------------------------
ethernet1/1/5   Desg 128.276 128  500  FWD  0    AUTO      No
ethernet1/1/6   Desg 128.280 128  500  LRN  0    AUTO      No
```

Loop guard

To configure loop guard on an interface:

- Enable loop guard on a per-port or port-channel interface in INTERFACE mode.

  ```
  spanning-tree guard {loop | root | none}
  ```

  - `loop` — Enables loop guard on an interface.
  - `root` — Enables root on an interface.
  - `none` — Enables the guard mode to none.
- Disable loop guard on a port or port-channel interface in INTERFACE mode.

  ```
  no spanning-tree guard loop
  ```

Port enabled with loop guard conditions

- Loop guard is supported on any STP-enabled port or port-channel interface in RPVST+ mode.
- You cannot enable root guard and loop guard at the same time on an STP port. The loop guard configuration overwrites an existing root guard configuration and vice versa.
- Enabling BPDU guard and loop guard at the same time on a port results in a port that remains in the Blocking state and prevents traffic from flowing through it. For example, when you configure both Portfast BPDU guard and loop guard:
  - If a BPDU is received from a remote device, BPDU guard places the port in the Err-Disabled Blocking state and no traffic forwards on the port.
  - If no BPDU is received from a remote device which was sending BPDUs, loop guard places the port in the Loop-Inconsistent Blocking state and no traffic forwards on the port.
- When used in a PVST+ network, STP loop guard performs per-port or per port-channel at a VLAN level. If no BPDUs are received on a port-channel interface, the port or port-channel transitions to a Loop-Inconsistent or Blocking state only for this VLAN.

Global parameters

All non-root bridges accept the timer values on the root bridge.

Forward-time	Amount of time required for an interface to transition from the Discarding state to the Learning state or from the Learning state to the Forwarding state.
Hello-time	Time interval within which the bridge sends BPDUs.
Max-age	Length of time the bridge maintains configuration information before it refreshes information by recomputing the RPVST+ topology.

- Modify the forward-time in seconds in CONFIGURATION mode, from 4 to 30, default 15.

  ```
  spanning-tree vlan vlan-id forward-time seconds
  ```
- Modify the hello-time in seconds in CONFIGURATION mode, from 1 to 10, default 2. With large configurations involving more numbers of ports, Dell EMC recommends increasing the hello-time.

  ```
  spanning-tree vlan vlan-id hello-time seconds
  ```

- Modify the max-age (in seconds) in CONFIGURATION mode, from 6 to 40, default 20.

```
spanning-tree vlan vlan-id max-age seconds
```

View RPVST+ global parameters

```
OS10# show spanning-tree active
Spanning tree enabled protocol rapid-pvst with force-version rstp
VLAN 1
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32769, Address 90b1.1cf4.a523
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 32769, Address 90b1.1cf4.a523
We are the root of VLAN 1
Configured hello time 2, max age 20, forward delay 15
```

Setting spanning-tree link type for rapid state transitions

To fasten the spanning-tree state transitions, you can set the link type to point-to-point. To set the link type to point-to-point:

- Use the following command in INTERFACE mode.

```
spanning-tree link-type point-to-point
```

RPVST+ commands

clear spanning-tree counters

Clears the counters for STP.

Syntax

```
clear spanning-tree counters [interface {ethernet node/slot/port[:subport] | port—channel number}}]
```

Parameters

- `interface` — Enter the interface type:
  - `ethernet node/slot/port[:subport]` — Deletes the spanning-tree counters from a physical port.
  - `port-channel number` — Deletes the spanning-tree counters for a port-channel interface, from 1 to 128.

Default Not configured

Command Mode EXEC

Usage Information Clear all STP counters on the device per the Ethernet interface or port-channel.

Example

```
OS10# clear spanning-tree counters interface port-channel 10
```

Supported Releases 10.2.0E or later

clear spanning-tree detected-protocol

Forces the MST ports to renegotiate with neighbors.

Syntax

```
clear spanning-tree detected-protocol [interface {ethernet node/slot/port[:subport] | port-channel number}]
```

Parameters

- `interface` — Enter the interface type:
  - `ethernet node/slot/port[:subport]` — Enter the Ethernet interface information, from 1 to 48.
  - `port-channel number` — Enter the port-channel number, from 1 to 128.

Default Not configured

Command Mode EXEC

Usage Information Use this command to force the RPVST+ port to re-negotiate with neighbors. If you use this command without parameters, the command applies to each device port.

Example

```
OS10# clear spanning-tree detected-protocol interface ethernet 1/1/1
```

Supported Release 10.2.0E or later

show spanning-tree vlan

Displays RPVST+ status and configuration information by VLAN ID.

Syntax `show spanning-tree vlan vlan-id`

Parameters `vlan vlan-id` — Enter the VLAN ID number, from 1 to 4093.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show spanning-tree
Spanning tree enabled protocol rapid-pvst
VLAN 1
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32769, Address 74e6.e2f5.bb80
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 32769, Address 74e6.e2f5.bb80
We are the root of VLAN 1
Configured hello time 2, max age 20, forward delay 15
Interface                                                  Designated
Name           PortID  Prio Cost     Sts Cost Bridge ID   PortID
-------------------------------------------------------------------------
ethernet1/1/1  128.260 128  200000000 FWD 0   32769   0000.0000.0000 128.260
ethernet1/1/2  128.264 128  200000000 FWD 0   32769   0000.0000.0000 128.264
ethernet1/1/3  128.268 128  200000000 FWD 0   32769   0000.0000.0000 128.268
ethernet1/1/4  128.272 128  200000000 FWD 0   32769   0000.0000.0000 128.272
ethernet1/1/5  128.276 128  200000000 FWD 0   32769   0000.0000.0000 128.276
ethernet1/1/6  128.280 128  200000000 FWD 0   32769   0000.0000.0000 128.280
ethernet1/1/7  128.284 128  200000000 FWD 0   32769   0000.0000.0000 128.284
ethernet1/1/8  128.288 128  200000000 FWD 0   32769   0000.0000.0000 128.288
ethernet1/1/9  128.292 128  200000000 FWD 0   32769   0000.0000.0000 128.292
ethernet1/1/10 128.296 128  200000000 FWD 0   32769   0000.0000.0000 128.296
ethernet1/1/11 128.300 128  200000000 FWD 0   32769   0000.0000.0000 128.300
ethernet1/1/12 128.304 128  200000000 FWD 0   32769   0000.0000.0000 128.304
```

Supported Releases 10.2.0E or later

spanning-tree bpdufilter

Enables or disables BPDU filtering on an interface.

Syntax	`spanning-tree bpdufilter {enable | disable}`
Parameters	• `enable` — Enables the BPDU filter on an interface. • `disable` — Disables the BPDU filter on an interface.
Default	Disabled
Command Mode	INTERFACE
Usage Information	Use the `enable` parameter to enable BPDU filtering.
Example	`OS10(conf-if-eth1/1/4)# spanning-tree bpdufilter enable`
Supported Releases	10.2.0E or later

spanning-tree bpduguard

Enables or disables the BPDU guard on an interface.

Syntax	`spanning-tree bpduguard {enable | disable}`
Parameters	• `enable` — Enables the BPDU guard filter on an interface. • `disable` — Disables the BPDU guard filter on an interface.
Default	Disabled
Command Mode	INTERFACE
Usage Information	BPDU guard prevents a port from receiving BPDUs. If the port receives a BPDU, it is placed in the Error-Disabled state.
Example	`OS10(conf-if-eth1/1/4)# spanning-tree bpduguard enable`
Supported Releases	10.2.0E or later

spanning-tree disable

Disables Spanning-Tree mode configured with the `spanning-tree mode` command globally on the switch or on specified interfaces.

Syntax	`spanning-tree disable`
Parameters	None
Default	Not configured.
Usage Information	The `no` version of this command re-enables STP and applies the currently configured spanning-tree settings.
Command Mode	CONFIGURATION INTERFACE

Example

```
OS10(config)# interface ethernet 1/1/4
OS10(config-if-eth1/1/4)# spanning-tree disable
```

Supported Releases 10.3.0E or later

spanning-tree guard

Enables or disables loop guard or root guard on an interface.

Syntax `spanning-tree guard {loop | root | none}`

Parameters

- `loop` — Enables loop guard on an interface.
- `root` — Enables root guard on an interface.
- `none` — Sets the guard mode to none.

Default Not configured

Usage Information Root guard and loop guard configurations are mutually exclusive. Configuring one overwrites the other from the active configuration.

Command Mode INTERFACE

Example

```
OS10(conf-if-eth1/1/4)# spanning-tree guard root
```

Supported Releases 10.2.0E or later

spanning-tree link-type

Sets the spanning-tree link type for faster convergence.

Syntax `spanning-tree link-type {auto | point-to-point | shared}`

Parameters

- `auto` — Enter the keyword to sets the link type based on the duplex setting of the interface.
- `point-to-point`—Specifies that the interface is a point-to-point or full-duplex link.
- `shared`—Specifies that the interface is a half-duplex medium.

Default Auto

Command Mode INTERFACE

Usage Information To enable fast state transition, use the `point-to-point` option.

Example

```
OS10(config)# spanning-tree link-type point-to-point
```

Supported Releases 10.4.3.0 or later

spanning-tree mode

Enables an STP type: RSTP, Rapid-PVST+, or MST.

Syntax `spanning-tree mode {rstp | mst | rapid-pvst}`

Parameters

- `rstp` — Sets STP mode to RSTP.
- `mst` — Sets STP mode to MST.
- `rapid-pvst` — Sets STP mode to RPVST+.

Default RPVST+

Command Mode CONFIGURATION

Usage Information All STP instances stop in the previous STP mode, and restart in the new mode. You can also change to RSTP/MST mode.

Example (RSTP)

```
OS10(config)# spanning-tree mode rstp
```

Example (MST)

```
OS10(config)# spanning-tree mode mst
```

Supported Releases 10.2.0E or later

spanning-tree port

Sets the port type as the EdgePort.

Syntax `spanning-tree port type edge`

Parameters None

Default Not configured

Command Mode INTERFACE

Usage Information When you configure an EdgePort on a device running STP, the port immediately transitions to the Forwarding state. Only configured ports connected to end hosts act as EdgePorts.

Example

```
OS10(config)# spanning-tree port type edge
```

Supported Releases 10.2.0E or later

spanning-tree vlan cost

Sets the path cost of the interface per VLAN for PVST calculations.

Syntax `spanning-tree vlan` *`vlan-id`* `cost {`*`value`*`}`

Parameters *`value`* — Enter a port cost value to set the path cost of the interface for PVST calculations, from 1 to 200000000.

Defaults

- 100- Mb/s Ethernet interface = 200000
- 1 Gigabit Ethernet interface = 20000
- 10-Gigabit Ethernet interface = 2000
- Port–channel interface with one 100 Mb/s Ethernet = 200000
- Port-channel interface with one 1 Gigabit Ethernet = 20000
- Port-channel interface with one 10 Gigabit Ethernet = 2000
- Port-channel with two 1 Gigabit Ethernet = 18000
- Port-channel with two 10 Gigabit Ethernet = 1800

- Port-channel with two 100 Mbps Ethernet = 180000

Command Mode	INTERFACE
Usage Information	The media speed of a LAN interface determines the STP port path cost default value.
Example	`OS10(conf-if-eth1/1/4)# spanning-tree vlan 10 cost 1000`
Supported Releases	10.2.0E or later

spanning-tree vlan disable

Disables spanning tree on a specified VLAN.

Syntax	`spanning-tree vlan` *`vlan-id`* `disable`
Parameters	*`vlan-id`* — Enter the VLAN ID number, from 1 to 4093.
Default	Enabled
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command enables spanning tree on the specified VLAN.
Example	`OS10(config)# spanning-tree vlan 100 disable`
Supported Releases	10.4.0E(R1) or later

spanning-tree vlan forward-time

Configures a time interval for the interface to wait in the Blocking state or Learning state before moving to the Forwarding state.

Syntax	`spanning-tree vlan` *`vlan-id`* `forward-time` *`seconds`*
Parameters	• *`vlan-id`*— Enter a VLAN ID number, from 1 to 4093. • *`seconds`* — Enter the forward-delay time in seconds, from 4 to 30.
Default	15 seconds
Command Mode	CONFIGURATION
Usage Information	None
Example	`OS10(config)# spanning-tree vlan 10 forward-time 16`
Supported Releases	10.2.0E or later

spanning-tree vlan force-version

Configures a forced version of spanning-tree to transmit BPDUs.

Syntax	`spanning-tree vlan` *`vlan-id`* `force-version {stp | rstp}`
Parameters	• `stp` — Forces the version for the BPDUs transmitted by RPVST+ to STP. • `rstp` — Forces the version for the BPDUs transmitted by RPVST+ to RSTP

Default	Not configured
Command Mode	CONFIGURATION
Usage Information	Forces a bridge that supports RPVST+ to operate in a STP-compatible mode.
Example	`OS10(config)# spanning-tree mst force-version`
Supported Releases	10.2.0E or later

spanning-tree vlan hello-time

Sets the time interval between generation and transmission of RPVST BPDUs.

Syntax	`spanning-tree vlan vlan-id hello-time seconds`
Parameters	• *vlan-id* — Enter the VLAN ID number, from 1 to 4093. • *seconds* — Enter a hello-time interval value in seconds, from 1 to 10.
Default	2 seconds
Command Mode	CONFIGURATION
Usage Information	Dell EMC recommends increasing the hello-time for large configurations, especially configurations with multiple ports.
Example	`OS10(config)# spanning-tree vlan 10 hello-time 5`
Supported Releases	10.2.0E or later

spanning-tree vlan mac-flush-threshold

Configures the MAC-flush threshold value for the specified VLAN.

Syntax	`spanning-tree vlan vlan-id mac-flush-threshold threshold-value`
Parameters	• *vlan-id* — Enter the spanning-tree VLAN ID number, from 1 to 4093. • *threshold-value*—Enter the threshold value for the number of flushes, from 0 to 65535. The default value is 0.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The threshold value indicates the number of port-based flush requests allowed to be invoked before starting the flush optimization. When the flush interval value is non-zero, port-and-instance-based flushing is triggered until the threshold is reached. Once the threshold is reached, VLAN-based flushing is triggered and the timer starts. The `no` form of the command resets the flush indication threshold of the specific instance to its default value.
Example	`OS10(config)# spanning-tree vlan 100 mac-flush-threshold 255`
Supported Releases	10.4.0E(R1) or later

spanning-tree vlan max-age

Configures the time period the bridge maintains configuration information before refreshing the information by recomputing RPVST.

Syntax	`spanning-tree vlan vlan-id max-age seconds`
Parameters	`max-age seconds` — Enter a maximum age value in seconds, from 6 to 40.
Default	20 seconds
Command Mode	CONFIGURATION
Usage Information	None
Example	`OS10(config)# spanning-tree vlan 10 max-age 10`
Supported Releases	10.2.0E or later

spanning-tree vlan priority

Sets the priority value for RPVST+.

Syntax	`spanning-tree vlan vlan-id priority priority value`
Parameters	`priority priority value` — Enter a bridge-priority value in increments of 4096, from 0 to 61440. Valid priority values are: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. All other values are rejected.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The RPVST+ protocol determines the root bridge. Assign one bridge a lower priority to increase the probability of it being the root bridge. A lower `priority value` increases the probability of the bridge becoming a root bridge.
Example	`OS10(config)# spanning-tree vlan 10 priority 0`
Supported Releases	10.2.0E or later

spanning-tree vlan priority (Interface)

Sets an interface priority when two bridges compete for position as the root bridge.

Syntax	`spanning-tree vlan vlan-id priority value`
Parameters	`value` — Enter a priority value in the increments of 16, from 0 to 240.
Default	128
Command Mode	INTERFACE
Usage Information	Breaks the tie between the two bridges which compete for root bridge.
Example	`OS10(conf-if-eth1/1/4)# spanning-tree vlan 10 priority 16`
Supported Releases	10.2.0E or later

spanning-tree vlan root

Designates a device as the primary or secondary root bridge.

Syntax	`spanning-tree vlan vlan-id root {primary | secondary}`
Parameters	• *vlan-id* — Enter a VLAN ID number, from 1 to 4093. • `root` — Designate the bridge as the primary or secondary root. • `primary` — Designate the bridge as the primary or root bridge. • `secondary` — Designate the bridge as the secondary or secondary root bridge.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	None
Example	`OS10(config)# spanning-tree vlan 1 root primary`
Supported Releases	10.2.0E or later

Rapid Spanning-Tree Protocol

Rapid Spanning-Tree Protocol (RSTP) is similar to STP, but provides faster convergence and interoperability with devices configured with STP and MSTP. RSTP is disabled by default. All enabled interfaces in L2 mode automatically add to the RSTP topology.

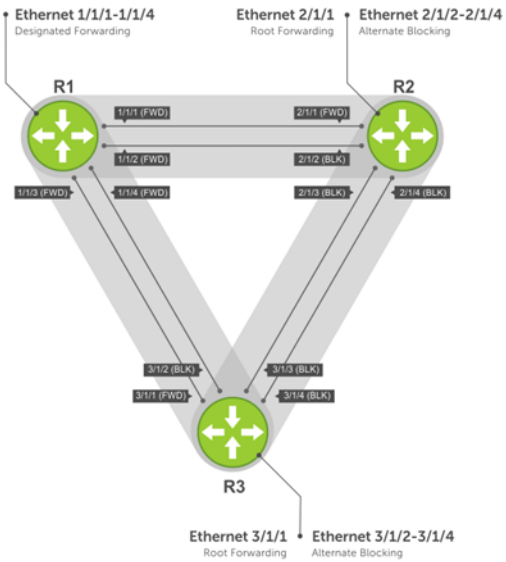
Configuring RSTP is a two-step process:

1 Ensure that the interfaces are in L2 mode.
2 Globally enable RSTP.

Enable globally

RSTP enables STP on all physical and port-channel interfaces which are in L2 mode to automatically include the interfaces as part of the RSTP topology. Only one path from a bridge to any other bridge is enabled. Bridges block a redundant path by disabling one of the link ports.

- Configure Spanning-Tree mode to RSTP in CONFIGURATION mode.
  ```
  spanning-tree mode rstp
  ```
- Disable RSTP globally for all L2 interfaces in CONFIGURATION mode.
  ```
  spanning-tree disable
  ```
- Remove an interface from the RSTP topology in INTERFACE mode.
  ```
  spanning-tree disable
  ```
- Re-enable an interface in INTERFACE mode.
  ```
  no spanning-tree disable
  ```
- Re-enable RSTP globally for all L2 interfaces in CONFIGURATION mode.
  ```
  no spanning-tree disable
  ```



View all port participating in RSTP

```
OS10# show spanning-tree
Spanning tree enabled protocol rstp with force-version rstp
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32768, Address 3417.4455.667f
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 32768, Address 90b1.1cf4.a523
Configured hello time 2, max age 20, forward delay 15
Interface                                                    Designated
Name               PortID    Prio  Cost      Sts   Cost Bridge ID            PortID
-------------------------------------------------------------------------------
ethernet1/1/1      128.260   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/2      128.264   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/3      128.268   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/4      128.272   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/5:1    128.276   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/5:2    128.277   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/5:3    128.278   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/5:4    128.279   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/6:1    128.280   128   2000      FWD   0    32768     3417.4455.667f  128.150
ethernet1/1/6:2    128.281   128   2000      FWD   0    32768     3417.4455.667f  128.151
ethernet1/1/6:3    128.282   128   2000      FWD   0    32768     3417.4455.667f  128.152
ethernet1/1/6:4    128.283   128   2000      BLK   0    32768     3417.4455.667f  128.153
ethernet1/1/7      128.284   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/8      128.288   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/9      128.292   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/10     128.296   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/11     128.300   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/12     128.304   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/13     128.308   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/14     128.312   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/15     128.316   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/16     128.320   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/17     128.324   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/18     128.328   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/19     128.332   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/20     128.336   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/21     128.340   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/22     128.344   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/23     128.348   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/24     128.352   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/25     128.356   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/26     128.360   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/27     128.364   128   200000000 BLK   0    0         0000.0000.0000  0.0
```

```
ethernet1/1/28    128.368   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/29    128.372   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/30    128.376   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/31    128.380   128   200000000 BLK   0    0         0000.0000.0000  0.0
ethernet1/1/32    128.384   128   200000000 BLK   0    0         0000.0000.0000  0.0
Interface
Name              Role  PortID    Prio Cost      Sts  Cost Link-type Edge
-------------------------------------------------------------------------
ethernet1/1/1     Disb  128.260   128  200000000 BLK  0    AUTO      No
ethernet1/1/2     Disb  128.264   128  200000000 BLK  0    AUTO      No
ethernet1/1/3     Disb  128.268   128  200000000 BLK  0    AUTO      No
ethernet1/1/4     Disb  128.272   128  200000000 BLK  0    AUTO      No
ethernet1/1/5:1   Disb  128.276   128  200000000 BLK  0    AUTO      No
```

Global parameters

The root bridge sets the values for forward-time, hello-time, and max-age, and overwrites the values set on other bridges participating in the RSTP group.

NOTE: Dell EMC recommends that only experienced network administrators change the RSTP group parameters. Poorly planned modification of the RSTP parameters can negatively affect network performance.

Forward-time	15 seconds — Amount of time an interface waits in the Listening state and Learning state before it transitions to the Forwarding state.
Hello-time	2 seconds — Time interval in which the bridge sends RSTP BPDUs.
Max-age	20 seconds — Length of time the bridge maintains configuration information before it refreshes that information by recomputing the RSTP topology.
Port cost	Port cost values to set the path cost of the interface: • 100-Mb/s Ethernet interfaces — 200000 • 1-Gigabit Ethernet interfaces — 20000 • 10-Gigabit Ethernet interfaces — 2000 • 40-Gigabit Ethernet interfaces — 500 • Port-channel with 100 Mb/s Ethernet interfaces — 200000 • Port-channel with 1-Gigabit Ethernet interfaces — 20000 • Port-channel with 10-Gigabit Ethernet interfaces — 2000 • Port-channel with 1x40Gigabit Ethernet interface — 500 • Port-channel with 2x40Gigabit Ethernet interfaces — 250

- Change the forward-time in CONFIGURATION mode, from 4 to 30, default 15.
  ```
  spanning-tree rstp forward-time seconds
  ```
- Change the hello-time in CONFIGURATION mode, from 1 to 10, default 2. With large configurations, especially those configurations with more ports, Dell EMC recommends increasing the hello-time.
  ```
  spanning-tree rstp hello-time seconds
  ```
- Change the max-age in CONFIGURATION mode, from 6 to 40, default 20.
  ```
  spanning-tree rstp max-age seconds
  ```

View current interface parameters

```
OS10# show spanning-tree active

Spanning tree enabled protocol rstp with force-version rstp
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32768, Address 90b1.1cf4.9b8a
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 32768, Address 90b1.1cf4.9b8a
We are the root
```

```
Configured hello time 2, max age 20, forward delay 15
Interface                                          Designated
Name          PortID  Prio Cost Sts Cost Bridge ID  PortID
-------------------------------------------------------------------
ethernet1/1/1 244.128 128  500  BLK 0    32768 90b1.1cf4.9b8a  128.244
ethernet1/1/2 248.128 128  500  BLK 0    32768 90b1.1cf4.9b8a  128.248
ethernet1/1/3 252.128 128  500  FWD 0    32768 90b1.1cf4.9b8a  128.252
ethernet1/1/4 256.128 128  500  BLK 0    32768 90b1.1cf4.9b8a  128.256
Interface
Name          Role PortID  Prio Cost Sts Cost Link-type Edge
-------------------------------------------------------------
ethernet1/1/1 Altr 128.244 128  500  BLK 0    AUTO      No
ethernet1/1/2 Altr 128.248 128  500  BLK 0    AUTO      No
ethernet1/1/3 Root 128.252 128  500  FWD 0    AUTO      No
ethernet1/1/4 Altr 128.256 128  500  BLK 0    AUTO      No
```

Interface parameters

Set the port cost and port priority values on interfaces in L2 mode.

Port cost	Value based on the interface type. The previous table lists the default values. The greater the port cost, the less likely the port is selected as a forwarding port.
Port priority	Influences the likelihood a port is selected to be a forwarding port in case several ports have the same port cost.

- Change the port cost of an interface in INTERFACE mode, from 1 to 200000000.

  ```
  spanning-tree rstp cost cost
  ```

- Change the port priority of an interface in INTERFACE mode, from 0 to 240, default 128.

  ```
  spanning-tree rstp priority priority-value
  ```

View current global parameter values

```
OS10# show spanning-tree active

Spanning tree enabled protocol rstp with force-version rstp
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32768, Address 90b1.1cf4.9b8a
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 32768, Address 90b1.1cf4.9b8a
We are the root
Configured hello time 2, max age 20, forward delay 15
Interface                                          Designated
Name          PortID  Prio Cost Sts Cost Bridge ID  PortID
-------------------------------------------------------------------
ethernet1/1/1 244.128 128  500  BLK 0    32768 90b1.1cf4.9b8a  128.244
ethernet1/1/2 248.128 128  500  BLK 0    32768 90b1.1cf4.9b8a  128.248
ethernet1/1/3 252.128 128  500  FWD 0    32768 90b1.1cf4.9b8a  128.252
ethernet1/1/4 256.128 128  500  BLK 0    32768 90b1.1cf4.9b8a  128.256
Interface
Name          Role PortID  Prio Cost Sts Cost Link-type Edge
-------------------------------------------------------------
ethernet1/1/1 Altr 128.244 128  500  BLK 0    AUTO      No
ethernet1/1/2 Altr 128.248 128  500  BLK 0    AUTO      No
ethernet1/1/3 Root 128.252 128  500  FWD 0    AUTO      No
ethernet1/1/4 Altr 128.256 128  500  BLK 0    AUTO      No
```

Root bridge selection

RSTP determines the root bridge. Assign one bridge a lower priority to increase the likelihood that it is selected as the root bridge.

- Assign a number as the bridge priority or designate it as the primary or secondary root bridge in CONFIGURATION mode. Configure the priority value range, from 0 to 65535 in multiples of 4096, default 32768. The lower the number assigned, the more likely the bridge becomes the root bridge.

  ```
  spanning-tree rstp priority priority-value
  ```

View bridge priority and root bridge assignment

```
OS10# show spanning-tree active
Spanning tree enabled protocol rstp with force-version rstp
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32768, Address 3417.4455.667f
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 36864, Address 90b1.1cf4.a523
Configured hello time 2, max age 20, forward delay 15
Interface                                                  Designated
Name              PortID   Prio  Cost  Sts  Cost Bridge ID          PortID
-------------------------------------------------------------------------
ethernet1/1/6:3   128.282  128   2000  FWD  0    32768  3417.4455.667f  128.152
ethernet1/1/6:4   128.283  128   2000  BLK  0    32768  3417.4455.667f  128.153
Interface
Name              Role  PortID  Prio Cost Sts Cost Link-type Edge
-----------------------------------------------------------------
ethernet1/1/6:3   Root  128.282 128  2000 FWD 0    AUTO      No
ethernet1/1/6:4   Altr  128.283 128  2000 BLK 0    AUTO      No
```

EdgePort forward traffic

EdgePort allows the interface to forward traffic approximately 30 seconds sooner as it skips the Blocking and Learning states. The `spanning-tree bpduguard enable` command causes the interface hardware to shut down when it receives a BPDU.

CAUTION: Configure EdgePort only on links connecting to an end station. EdgePort can cause loops if you enable it on an interface connected to a network.

- Enable EdgePort on an interface in INTERFACE mode.

  ```
  spanning-tree port type edge
  ```

Configure EdgePort and view interface status

```
OS10(conf-if-eth1/1/5)# spanning-tree port type edge
```

View interface status

```
OS10# show spanning-tree interface ethernet 1/1/5
ethernet1/1/5 of RSTP 1 is designated Forwarding
Edge port:yes port guard :none (default)
Link type is point-to-point (auto)
Boundary: YES bpdu filter :disable bpdu guard :disable  bpduguard shutdown-on-
violation :disable  RootGuard:  disable LoopGuard  disable
Bpdus (MRecords) sent 610, received 5
Interface                                              Designated
Name            PortID   Prio Cost Sts  Cost Bridge ID       PortID
-----------------------------------------------------------------------
ethernet1/1/5   128.272  128  500  FWD  0    32768 90b1.1cf4.a911  128.272
=======================================================================
```

Spanning-tree extensions

STP extensions ensure efficient network convergence by securely enforcing the active network topology. OS10 supports BPDU filtering, BPDU guard, loop guard, and root guard STP extensions.

BPDU filtering	Protects the network from unexpected flooding of BPDUs from an erroneous device. Enabling BPDU Filtering instructs the hardware to drop BPDUs and prevents flooding from reaching the CPU. BPDU filtering is enabled by default on Edge ports. All BPDUs received on the Edge port drop. If you explicitly configure BPDU filtering on a port, that port drops all BPDUs it receives.
BPDU guard	Blocks the L2 bridged ports and LAG ports connected to end hosts and servers from receiving any BPDUs. When you enable BPDU guard, it places a bridge or LAG port in an Error_Disable or Blocking state if the port receives any BPDU frames. In a LAG, all member ports, including new members are placed in an Blocking state. The network traffic drops but the port continues to forward BPDUs to the CPU that are later dropped. To prevent further reception of BPDUs, configure a port to shut down using the `shutdown` command. The port can only resume operation from the Shutdown state after manual intervention.
Root guard	Avoids bridging loops and preserves the root bridge position during network transitions. STP selects the root bridge with the lowest priority value. During network transitions, another bridge with a lower priority may attempt to become the root bridge and cause unpredictable network behavior. To avoid such an attempt and preserve the position of the root bridge, configure the `spanning-tree guard root` command. Root guard is enabled on ports that are designated ports. The root guard configuration applies to all VLANs configured on the port.
Loop guard	Prevents L2 forwarding loops caused by a cable or interface hardware failure. When a hardware failure occurs, a participating spanning tree link becomes unidirectional and a port stops receiving BPDUs. When a blocked port stops receiving BPDUs, it transitions to a Forwarding state causing spanning tree loops in the network. You can enable loop guard on a port that transitions to the Loop-Inconsistent state until it receives BPDUs using the `spanning-tree guard loop` command. After BPDUs are received, the port moves out of the Loop-Inconsistent or blocking state and transitions to an appropriate state determined by STP. Enabling loop guard on a per port basis enables it on all VLANs configured on the port. If you disable loop guard on a port, it is moved to the Listening state.

If you enable BPDU Filter and BPDU Guard on the same port, the BPDU Filter configuration takes precedence. Root Guard and Loop Guard are mutually exclusive. Configuring one overwrites the other from the active configuration.

- Enable spanning-tree BPDU filter in INTERFACE mode. Use the `spanning-tree bpdufilter disable` command to disable the BPDU filter on the interface.

  ```
  spanning-tree bpdufilter enable
  ```

- Enable spanning-tree BPDU guard in INTERFACE mode.

  ```
  spanning-tree bpduguard enable
  ```

  - Use the `shutdown` command to shut down the port channel interface, all member ports that are disabled in the hardware.
  - Use the `spanning-tree bpduguard disable` command to add a physical port to a port-channel already in the Error Disable state, the new member port is also disabled in the hardware.

- Set the guard types to avoid loops in INTERFACE mode.

  ```
  spanning-tree guard {loop | root | none}
  ```

  - `loop` — Set the guard type to loop.
  - `none` — Set the guard type to none.
  - `root` — Set the guard type to root.

BPDU filter

```
OS10(conf-if-eth1/1/4)# spanning-tree bpdufilter enable
OS10(conf-if-eth1/1/4)# do show spanning-tree interface ethernet 1/1/4
ethernet1/1/4 of vlan1 is designated Blocking
Edge port:no (default) port guard :none (default)
Link type is point-to-point (auto)
Boundary: NO  bpdu filter : Enable bpdu guard :  bpduguard shutdown-on-
```

```
violation :disable  RootGuard:  enable LoopGuard  disable
Bpdus (MRecords) sent 134, received 138
Interface                                                  Designated
Name           PortID   Prio Cost Sts  Cost  Bridge ID          PortID
-------------------------------------------------------------------------
ethernet1/1/4  128.272  128  500  BLK  500   32769    90b1.1cf4.a911 128.272
```

BPDU guard

```
OS10(config)# interface ethernet 1/1/4
OS10(conf-if-eth1/1/4)# spanning-tree bpduguard enable
OS10(conf-if-eth1/1/4)# exit
OS10(config)# interface ethernet 1/1/4
OS10(conf-if-eth1/1/4)# do show spanning-tree interface ethernet 1/1/4
ethernet1/1/4 of vlan1 is designated Blocking
Edge port:no (default) port guard :none (default)
Link type is point-to-point (auto)
Boundary: NO  bpdu filter : Enable bpdu guard :  bpduguard shutdown-on-
violation :enable  RootGuard:  enable LoopGuard  disable
Bpdus (MRecords) sent 134, received 138
Interface                                                 Designated
Name           PortID  Prio Cost Sts  Cost Bridge ID         PortID
-----------------------------------------------------------------
ethernet1/1/4  128.272 128  500  BLK  500  32769    90b1.1cf4.a911 128.272
```

Loop guard

```
OS10(config)# interface ethernet 1/1/4
OS10(conf-if-eth1/1/4)# spanning-tree guard loop
OS10(conf-if-eth1/1/4)# do show spanning-tree interface ethernet 1/1/4
ethernet1/1/4 of vlan1 is root Forwarding
Edge port:no (default) port guard :none (default)
Link type is point-to-point (auto)
Boundary: NO  bpdu filter : bpdu guard :  bpduguard shutdown-on-
violation :disable  RootGuard:  disable LoopGuard  enable
Bpdus (MRecords) sent 7, received 20
Interface                                                   Designated
Name           PortID  Prio  Cost Sts  Cost Bridge ID           PortID
-----------------------------------------------------------------------
ethernet1/1/4  128.272 128   500  FWD  0    32769    90b1.1cf4.9d3b 128.272
```

Root guard

```
OS10(conf-if-eth1/1/4)# spanning-tree guard root
OS10(conf-if-eth1/1/4)# do show spanning-tree interface ethernet 1/1/4
ethernet1/1/4 of vlan1 is root Forwarding
Edge port:no (default) port guard :none (default)
Link type is point-to-point (auto)
Boundary: NO  bpdu filter : bpdu guard :  bpduguard shutdown-on-
violation :disable  RootGuard:  enable LoopGuard  disable
Bpdus (MRecords) sent 7, received 33
Interface                                                   Designated
Name           PortID  Prio Cost Sts Cost Bridge ID        PortID
-----------------------------------------------------------------------
ethernet1/1/4  128.272 128  500  BLK 500  32769  90b1.1cf4.a911 128.272
```

Setting spanning-tree link type for rapid state transitions

To fasten the spanning-tree state transitions, you can set the link type to point-to-point. To set the link type to point-to-point:

- Use the following command in INTERFACE mode.

  ```
  spanning-tree link-type point-to-point
  ```

RSTP commands

clear spanning-tree counters

Clears the counters for STP.

Syntax

```
clear spanning-tree counters [interface {ethernet node/slot/port[:subport] | port—channel number}}]
```

Parameters

- `interface` — Enter the interface type:
  - `ethernet node/slot/port[:subport]` — Deletes the spanning-tree counters from a physical port.
  - `port-channel number` — Deletes the spanning-tree counters for a port-channel interface, from 1 to 128.

Default Not configured

Command Mode EXEC

Usage Information Clear all STP counters on the device per the Ethernet interface or port-channel.

Example

```
OS10# clear spanning-tree counters interface port-channel 10
```

Supported Releases 10.2.0E or later

show spanning-tree active

Displays the RSTP configuration and information for RSTP-active interfaces.

Syntax

```
show spanning-tree active
```

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show spanning-tree active

Spanning tree enabled protocol rstp with force-version rstp
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32768, Address 90b1.1cf4.9b8a
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 32768, Address 90b1.1cf4.9b8a
We are the root
Configured hello time 2, max age 20, forward delay 15
Interface                                        Designated
Name          PortID  Prio Cost Sts Cost Bridge ID   PortID
-------------------------------------------------------------------------
ethernet1/1/1 244.128 128  500  BLK 0    32768 90b1.1cf4.9b8a  128.244
ethernet1/1/2 248.128 128  500  BLK 0    32768 90b1.1cf4.9b8a  128.248
ethernet1/1/3 252.128 128  500  FWD 0    32768 90b1.1cf4.9b8a  128.252
ethernet1/1/4 256.128 128  500  BLK 0    32768 90b1.1cf4.9b8a  128.256
Interface
Name          Role PortID  Prio Cost Sts Cost Link-type Edge
```

```
-------------------------------------------------------------
ethernet1/1/1 Altr 128.244 128  500  BLK 0    AUTO      No
ethernet1/1/2 Altr 128.248 128  500  BLK 0    AUTO      No
ethernet1/1/3 Root 128.252 128  500  FWD 0    AUTO      No
ethernet1/1/4 Altr 128.256 128  500  BLK 0    AUTO      No
```

Supported Releases 10.2.0E or later

show spanning-tree interface

Displays spanning-tree interface information for Ethernet and port-channels.

Syntax `show spanning-tree interface {ethernet node/slot/port [:subport] | port-channel port-id} [detail]`

Parameters

- `ethernet node/slot/port[:subport]` — Displays spanning-tree information for a physical interface.
- `port-channel port-id` — Displays spanning-tree information for a port-channel number, from 1 to 128.
- `detail` — (Optional) Displays detailed information on the interface.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show spanning-tree interface ethernet 1/1/6:2 detail
Port 281 (ethernet1/1/6:2) of RSTP 1 is root Forwarding
Port path cost 2000, Port priority 128, Port Identifier 281.128
Designated root has priority 32768, address 34:17:44:55:66:7f
Designated bridge has priority 32768, address 34:17:44:55:66:7f
Designated port id is 151.128, designated path cost
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state 1
Link type is point-to-point by default, auto
PVST Simulation is enabled by default
BPDU sent 3, received 7
```

Supported Releases 10.2.0E or later

spanning-tree bpdufilter

Enables or disables BPDU filtering on an interface.

Syntax `spanning-tree bpdufilter {enable | disable}`

Parameters

- `enable` — Enables the BPDU filter on an interface.
- `disable` — Disables the BPDU filter on an interface.

Default Disabled

Command Mode INTERFACE

Usage Information Use the `enable` parameter to enable BPDU filtering.

Example

```
OS10(conf-if-eth1/1/4)# spanning-tree bpdufilter enable
```

Supported Releases 10.2.0E or later

spanning-tree bpduguard

Enables or disables the BPDU guard on an interface.

Syntax `spanning-tree bpduguard {enable | disable}`

Parameters

- `enable` — Enables the BPDU guard filter on an interface.
- `disable` — Disables the BPDU guard filter on an interface.

Default Disabled

Command Mode INTERFACE

Usage Information BPDU guard prevents a port from receiving BPDUs. If the port receives a BPDU, it is placed in the Error-Disabled state.

Example

```
OS10(conf-if-eth1/1/4)# spanning-tree bpduguard enable
```

Supported Releases 10.2.0E or later

spanning-tree disable

Disables Spanning-Tree mode configured with the `spanning-tree mode` command globally on the switch or on specified interfaces.

Syntax `spanning-tree disable`

Parameters None

Default Not configured.

Usage Information The `no` version of this command re-enables STP and applies the currently configured spanning-tree settings.

Command Mode CONFIGURATION

INTERFACE

Example

```
OS10(config)# interface ethernet 1/1/4
OS10(config-if-eth1/1/4)# spanning-tree disable
```

Supported Releases 10.3.0E or later

spanning-tree guard

Enables or disables loop guard or root guard on an interface.

Syntax `spanning-tree guard {loop | root | none}`

Parameters

- `loop` — Enables loop guard on an interface.
- `root` — Enables root guard on an interface.
- `none` — Sets the guard mode to none.

Default Not configured

Usage Information	Root guard and loop guard configurations are mutually exclusive. Configuring one overwrites the other from the active configuration.
Command Mode	INTERFACE
Example	`OS10(conf-if-eth1/1/4)# spanning-tree guard root`
Supported Releases	10.2.0E or later

spanning-tree link-type

Sets the spanning-tree link type for faster convergence.

Syntax	`spanning-tree link-type {auto | point-to-point | shared}`
Parameters	• `auto` — Enter the keyword to sets the link type based on the duplex setting of the interface. • `point-to-point`—Specifies that the interface is a point-to-point or full-duplex link. • `shared`—Specifies that the interface is a half-duplex medium.
Default	Auto
Command Mode	INTERFACE
Usage Information	To enable fast state transition, use the `point-to-point` option.
Example	`OS10(config)# spanning-tree link-type point-to-point`
Supported Releases	10.4.3.0 or later

spanning-tree mode

Enables an STP type: RSTP, Rapid-PVST+, or MST.

Syntax	`spanning-tree mode {rstp | mst | rapid-pvst}`
Parameters	• `rstp` — Sets STP mode to RSTP. • `mst` — Sets STP mode to MST. • `rapid-pvst` — Sets STP mode to RPVST+.
Default	RPVST+
Command Mode	CONFIGURATION
Usage Information	All STP instances stop in the previous STP mode, and restart in the new mode. You can also change to RSTP/MST mode.
Example (RSTP)	`OS10(config)# spanning-tree mode rstp`
Example (MST)	`OS10(config)# spanning-tree mode mst`
Supported Releases	10.2.0E or later

spanning-tree port

Sets the port type as the EdgePort.

Syntax	`spanning-tree port type edge`
Parameters	None
Default	Not configured
Command Mode	INTERFACE
Usage Information	When you configure an EdgePort on a device running STP, the port immediately transitions to the Forwarding state. Only configured ports connected to end hosts act as EdgePorts.
Example	`OS10(config)# spanning-tree port type edge`
Supported Releases	10.2.0E or later

spanning-tree rstp force-version

Configures a forced version of spanning tree to transmit BPDUs.

Syntax	`spanning-tree rstp force-version` *stp*
Parameters	*stp* — Force the version for the BPDUs transmitted by RSTP.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	Forces a bridge that supports RSTP or MST to operate in a STP-compatible manner to avoid frame misordering and duplication in known LAN protocols that are sensitive.
Example	`OS10(config)# spanning-tree rstp force-version stp`
Supported Releases	10.2.0E or later

spanning-tree rstp forward-time

Configures a time interval for the interface to wait in the Blocking state or Learning state before moving to the Forwarding state.

Syntax	`spanning-tree rstp forward-time` *seconds*
Parameters	*seconds* — Enter the number of seconds an interface waits in the Blocking or Learning States before moving to the Forwarding state, from 4 to 30.
Default	15 seconds
Command Mode	CONFIGURATION
Usage Information	None
Example	`OS10(config)# spanning-tree rstp forward-time 16`
Supported Releases	10.2.0E or later

spanning-tree rstp hello-time

Sets the time interval between generation and transmission of RSTP BPDUs.

Syntax	`spanning-tree rstp hello-time` *seconds*
Parameters	*seconds* — Enter a hello-time interval value in seconds, from 1 to 10.
Default	2 seconds
Command Mode	CONFIGURATION
Usage Information	Dell EMC recommends increasing the hello-time for large configurations, especially configurations with multiple ports.
Example	`OS10(config)# spanning-tree rstp hello-time 5`
Supported Releases	10.2.0E or later

spanning-tree rstp mac-opt-flush-threshold

Sets the flush indication threshold value on the RSTP instance.

Syntax	`spanning-tree rstp mac-opt-flush-threshold` *threshold-value*
Parameters	*threshold-value*—Enter the threshold value for the number of flushes, from 0 to 65535. The default value is 0.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The threshold value indicates the number of flush indications to go before the flush interval timer is triggered. When flush indication threshold is set to the default value and the flush interval is set to a non-default value, flushing occurs during the first flush indication trigger. When the flush indication threshold value is non-default (n) and flush interval value is non-default, port-based flushing is triggered until the threshold (n) is reached. Once the threshold is reached, flushing is triggered and the timer starts. The `no` form of the command sets the flush indication threshold to its default value.
Example	`OS10(config)# spanning-tree rstp mac-opt-flush-threshold 255`
Supported Releases	10.4.0E(R1) or later

spanning-tree rstp max-age

Configures the time period the bridge maintains configuration information before refreshing the information by recomputing the RSTP topology.

Syntax	`max-age` *seconds*
Parameters	*seconds* — Enter a maximum age value in seconds, from 6 to 40.
Default	20 seconds
Command Mode	CONFIGURATION
Usage Information	None

Example	`OS10(config)# spanning-tree rstp max-age 10`
Supported Releases	10.2.0E or later

spanning-tree rstp

Sets the priority value for RSTP.

Syntax	`spanning-tree rspt priority` *`priority value`*
Parameters	`priority` *`priority value`* — Enter a bridge-priority value in increments of 4096, from 0 to 61440. Valid priority values are: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. All other values are rejected.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	RSTP determines the root bridge but you can assign one bridge a lower priority to increase the probability of it being the root bridge. A lower *`priority value`* increases the probability of the bridge becoming a root bridge.
Example	`OS10(config)# spanning-tree rstp priority 200`
Supported Releases	10.2.0E or later

Virtual LANs

VLANs segment a single flat L2 broadcast domain into multiple logical L2 networks. Each VLAN is uniquely identified by a VLAN ID or tag consisting of 12 bits in the Ethernet frame. VLAN IDs range from 1 to 4093 and provide a total of 4093 logical networks.

You can assign ports on a single physical device to one or more VLANs creating multiple logical instances on a single physical device. The virtual logical switches spanning across different physical devices emulate multiple logically segmented L2 networks on a single physical network.

Each VLAN has its own broadcast domain. The unicast, multicast, and broadcast network traffic from ports that belong to a VLAN forwards or floods to ports in the same VLAN only. Traffic between VLANs routes from one VLAN to another. You can also assign each VLAN an IP address to group all the ports within a single IP subnet.

Segment a L2 network using VLANs to:

- Minimize broadcast and multicast traffic in the L2 network
- Increase security by isolating ports into different VLANs
- Ease network management

Default VLAN

All interface ports are administratively up in L2 mode and are automatically placed in the default VLAN as untagged interfaces.

When you assign a port to a non-default VLAN in Trunk mode, the interface remains an untagged member of the default VLAN and a tagged member of the new VLAN. When you assign a port to a non-default VLAN in Access mode, it removes from the default VLAN and is assigned to the new VLAN as an untagged member of the new VLAN.

- VLAN 1 is the default VLAN.
- You cannot delete the default VLAN. However, you can change the default VLAN ID number using the `default vlan-id` command.

Use the `show vlan` command to verify that the interface is part of the default VLAN (VLAN 1).

Default VLAN configuration

```
OS10# show vlan

Codes: * - Default VLAN, G-GVRP VLANs, R-Remote Port Mirroring VLANs, P-Primary, C-Community, I-
Isolated
Q: A-Access (Untagged), T-Tagged
   x-Dot1x untagged, X-Dot1x tagged
   G-GVRP tagged, M-Vlan-stack, H-VSN tagged
   i-Internal untagged, I-Internal tagged, v-VLT untagged, V-VLT tagged
    NUM    Status    Description   Q Ports
*   1      up                      A Eth1/1/1-1/1/54
```

Create or remove VLANs

You can create VLANs and add physical interfaces or port-channel LAG interfaces to the VLAN as tagged or untagged members. You can add an Ethernet interface as a trunk port or as an access port, but it cannot be added as both at the same time.

Multiple non-default vlans with physical and port channel ports in Access and Trunk modes

```
OS10# show vlan

Codes: * - Default VLAN, G-GVRP VLANs, R-Remote Port Mirroring VLANs, P-Primary, C-Community, I-
Isolated
Q: A-Access (Untagged), T-Tagged
   x-Dot1x untagged, X-Dot1x tagged
   G-GVRP tagged, M-Vlan-stack, H-VSN tagged
   i-Internal untagged, I-Internal tagged, v-VLT untagged, V-VLT tagged
    NUM    Status    Description                    Q Ports
*   1      up                                       A Eth1/1/2 1/1/3:2 1/1/3:3 1/1/3:4 1/1/4
1/1/5 1/1/6 1/1/7 1/1/8 1/1/9 1/1/10 1/1/11 1/1/12 1/1/13 1/1/14 1/1/15 1/1/16 1/1/17 1/1/18
1/1/19 1/1/20 1/1/21 1/1/22 1/1/23 1/1/24 1/1/25:1 1/1/25:2 1/1/25:3 1/1/25:4 1/1/26 1/1/27
1/1/28 1/1/30 1/1/32
                                                    A Po40
    200    up                                       T Eth1/1/3:2
                                                    T Po40
                                                    A Eth1/1/31
    320    up                                       T Eth1/1/25:4 1/1/32
                                                    T Po40
                                                    A Eth1/1/3:1
49 1/1/50 1/1/51 1/1/52 1/1/53 1/1/54
```

The `shutdown` command stops L3–routed traffic only. L2 traffic continues to pass through the VLAN. If the VLAN is not a routed VLAN configured with an IP address, the `shutdown` command has no effect on VLAN traffic.

When you delete a VLAN using the `no interface vlan` *vlan-id* command, any interfaces assigned to that VLAN are assigned to the default VLAN as untagged interfaces.

To configure a port-based VLAN, enter INTERFACE-VLAN mode for VLAN-related configuration tasks and create a VLAN. To enable the VLAN, assign member interfaces in L2 mode.

1 Create a VLAN and enter the VLAN number in INTERFACE mode, from 1 to 4093.
   `interface vlan` *vlan-id*
2 Delete a VLAN in CONFIGURATION mode.
   `no interface vlan` *vlan-id*

Create VLAN

```
OS10(config)# interface vlan 108
```

Delete VLAN

```
OS10(config)# no interface vlan 108
```

View configured VLANs

```
OS10(config)# do show interface vlan

Vlan 1 is up, line protocol is up
Address is , Current address is
Interface index is 69208865
Internet address is not set
MTU 1532 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 240
Last clearing of "show interface" counters Queueing strategy: fifo Time since last interface
status change:


Vlan 200 is up, line protocol is up
Address is , Current address is
Interface index is 69209064
Internet address is not set
MTU 1532 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 240
Last clearing of "show interface" counters Queueing strategy: fifo Time since last interface
status change:


Vlan 320 is up, line protocol is up
Address is , Current address is
Interface index is 69209184
Internet address is not set
MTU 1532 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 240
Last clearing of "show interface" counters Queueing strategy: fifo Time since last interface
status change:
```

Access mode

An access port is an untagged member of only one VLAN. Configure a port in Access mode and configure which VLAN carries the traffic for that interface. If you do not configure the VLAN for a port in Access mode, or an access port, the interface carries traffic for VLAN 1, the default VLAN.

Change the access port membership in a VLAN by specifying the new VLAN. You must create the VLAN before you can assign the port in Access mode to that VLAN. Use the `no switchport access vlan` command to reset to default VLAN.

1 Configure a port in INTERFACE mode.

```
interface ethernet node/slot/port[:subport]
```

2 Set the interface to Switchport mode as access in INTERFACE mode.

```
switchport mode access
```

3 Enter the VLAN number for the untagged port in INTERFACE mode.

```
switchport access vlan vlan-id
```

Configure port in Access mode

```
OS10(config)# interface ethernet 1/1/9
OS10(config-if-eth1/1/9)# switchport mode access
OS10(config-if-eth1/1/9)# switchport access vlan 604
```

Show running configuration

```
OS10# show running-configuration
...
!
 interface ethernet1/1/5
...
 switchport access vlan 604
 no shutdown
!
interface vlan1
 no shutdown
...
```

Trunk mode

A trunk port can be a member of multiple VLANs set up on an interface. A trunk port transmits traffic for all VLANs. To transmit traffic on a trunk port with multiple VLANs, OS10 uses tagging or the 802.1q encapsulation method.

1 Configure a port in INTERFACE mode.

```
interface ethernet node/slot/port[:subport]
```

2 Change Switchport mode to Trunk mode in INTERFACE mode.

```
switchport mode trunk
```

3 Enter the allowed VLANs on the trunk port in INTERFACE mode.

```
switchport trunk allowed vlan vlan-id
```

Configure port in Trunk mode

```
OS10(config)# interface ethernet 1/1/6
OS10(conf-if-eth1/1/6)# switchport mode trunk
OS10(conf-if-eth1/1/6)# switchport trunk allowed vlan 108
```

View running configuration

```
OS10# show running-configuration
...
!
 interface ethernet1/1/8
  switchport mode trunk
 switchport trunk allowed vlan 108
 no shutdown
!
interface vlan1
 no shutdown
!
...
```

Assign IP address

You can assign an IP address to each VLAN to make it a L3 VLAN. All the ports in that VLAN belong to that particular IP subnet.

The traffic between the ports in different VLANs route using the IP address. Configure the L3 VLAN interface to remain administratively UP or DOWN using the `shutdown` and `no shutdown` commands. This provisioning only affects the L3 traffic across the members of a VLAN and does not affect the L2 traffic.

You cannot assign an IP address to the default VLAN (VLAN 1). You can place VLANs and other logical interfaces in L3 mode to receive and send routed traffic.

1 Create a VLAN in CONFIGURATION mode, from 1 to 4093.

```
interface vlan vlan-id
```

2 Assign an IP address and mask to the VLAN in INTERFACE-VLAN mode.

```
ip address ip-address/prefix-length [secondary]
```

- *ip-address/prefix–length* — Enter the IP address in dotted-decimal A.B.C.D/x format.
- `secondary` — Enter the interface backup IP address.

Assign IP address to VLAN

```
OS10(config)# interface vlan 200
OS10(conf-if-vl-200)# ip address 10.1.15.1/8
```

View VLAN configuration

```
OS10(conf-if-vl-200)# do show interface vlan

Vlan 1 is up, line protocol is up
Address is , Current address is
Interface index is 69208865
Internet address is not set
MTU 1532 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 240
Last clearing of "show interface" counters Queueing strategy: fifo Time since last interface
status change:


Vlan 200 is up, line protocol is up
Address is , Current address is
Interface index is 69209064
Internet address is not set
MTU 1532 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 240
Last clearing of "show interface" counters Queueing strategy: fifo Time since last interface
status change:


Vlan 320 is up, line protocol is up
Address is , Current address is
Interface index is 69209184
Internet address is 20.2.11.1/24
MTU 1532 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 240
Last clearing of "show interface" counters Queueing strategy: fifo Time since last interface
status change:
```

View VLAN configuration

You can view configuration information related to VLANs using `show` commands.

- View the VLAN status and configuration information in EXEC mode.

```
show vlan
```

- View the VLAN interface configuration in EXEC mode.

```
show interfaces vlan
```

- View the VLAN interface configuration for a specific VLAN ID in EXEC mode.

```
show interfaces vlan vlan-id
```

View VLAN configuration

```
OS10# show vlan

Codes: * - Default VLAN, G-GVRP VLANs, R-Remote Port Mirroring VLANs, P-Primary, C-Community, I-
Isolated
Q: A-Access (Untagged), T-Tagged
   x-Dot1x untagged, X-Dot1x tagged
   G-GVRP tagged, M-Vlan-stack, H-VSN tagged
   i-Internal untagged, I-Internal tagged, v-VLT untagged, V-VLT tagged
    NUM    Status    Description       Q Ports
*   1      up                           A Eth1/1/1-1/1/32
                                        A Po40
    200    up                           T Eth1/1/3:2
                                        T Po40
                                        A Eth1/1/31
    320    up                           T Eth1/1/25:4 1/1/32
                                        T Po40
                                        A Eth1/1/3:1
```

View interface VLAN configuration

```
OS10# show interface vlan
Vlan 1 is up, line protocol is up
Address is , Current address is
Interface index is 69208865
Internet address is not set
MTU 1532 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 240
Last clearing of "show interface" counters Queueing strategy: fifo Time since last interface
status change:

Vlan 200 is up, line protocol is up
Address is , Current address is
Interface index is 69209064
Internet address is not set
MTU 1532 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 240
Last clearing of "show interface" counters Queueing strategy: fifo Time since last interface
status change:

Vlan 320 is up, line protocol is up
Address is , Current address is
Interface index is 69209184
Internet address is not set
MTU 1532 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 240
Last clearing of "show interface" counters Queueing strategy: fifo Time since last interface
status change:
```

View interface configuration for specific VLAN

```
OS10# show interface vlan 320
Vlan 320 is up, line protocol is up
Address is , Current address is
Interface index is 69209184
Internet address is not set
MTU 1532 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 240
Last clearing of "show interface" counters Queueing strategy: fifo Time since last interface
status change:
```

VLAN commands

description (VLAN)

Adds a description to the selected VLAN.

Syntax	`description description`
Parameters	*description* — Enter a text string to identify the VLAN. A maximum of 80 characters.
Default	Not configured
Command Mode	INTERFACE-VLAN
Usage Information	None
Example	`OS10(conf-if-vlan)# description vlan3`
Supported Releases	10.2.0E or later

interface vlan

Creates a VLAN interface.

Syntax	`interface vlan vlan-id`
Parameters	*vlan-id* — Enter the VLAN ID number, from 1 to 4093.
Default	VLAN 1
Command Mode	CONFIGURATION
Usage Information	FTP, TFTP, MAC ACLs, and SNMP operations are not supported. IP ACLs are supported on VLANs only. The `no` version of this command deletes the interface.
Example	`OS10(config)# interface vlan 10` `OS10(conf-if-vl-10)#`
Supported Releases	10.2.0E or later

show vlan

Displays VLAN configurations.

Syntax	`show vlan vlan-id`
Parameters	*vlan-id* — (Optional) Enter a VLAN ID number, from 1 to 4093.
Default	Not configured
Command Mode	EXEC
Usage Information	Use this command to view VLAN configuration information for a specific VLAN ID.

Example	`OS10(config)# do show vlan`
	`Codes: * - Default VLAN, M - Management VLAN, R - Remote Port Mirroring VLANs`
	`Q: A - Access (Untagged), T - Tagged`

```
    NUM    Status    Description                    Q Ports
*   1      Active                                   A Eth1/1/15
                                                    A Po100
    2101   Active                                   T Eth1/1/1,1/1/3
                                                    T Po100
    2102   Active                                   T Eth1/1/1,1/1/3
```

Supported Releases 10.2.0E or later

Port monitoring

Port monitoring monitors ingress or egress traffic of one port to another for analysis. A monitoring port (MG) or destination port, is the port where the monitored traffic is sent for analysis. A monitored port (MD) or source port is the source interface that is monitored for traffic analysis.

The different types of port monitoring are:

- **Local port monitoring** — Port monitoring is done in the same switch. The switch forwards a copy of incoming and outgoing traffic from one port to another port for further analysis.
- **Remote port monitoring (RPM)** — Port monitoring is done on traffic running across a remote device in the same network. The L2 network carries the monitored traffic.
- **Encapsulated remote port monitoring (ERPM)** — Port monitoring is done on the L3 network. The traffic from the source port is encapsulated and forwards to the destination port in another switch.

Local port monitoring

For local port monitoring, the monitored source ports and monitoring destination ports are on the same device.

Configure local monitoring session

1 Verify that the intended monitoring port has no configuration other than `no shutdown` and `no switchport`.

```
show running-configuration
```

2 Create a monitoring session in CONFIGURATION mode.

```
monitor session session-id [local]
```

3 Enter the source and direction of the monitored traffic in MONITOR-SESSION mode.

```
source interface interface-type {both | rx | tx}
```

4 Enter the destination of traffic in MONITOR-SESSION mode.

```
destination interface interface-type
```

Create monitoring session

```
OS10(config)# monitor session 1
OS10(conf-mon-local-1)#
```

Configure source and destination port, and traffic direction

```
OS10(conf-mon-local-1)# source interface ethernet 1/1/7-1/1/8 rx
OS10(conf-mon-local-1)# destination interface ethernet1/1/1
OS10(conf-mon-local-1)# no shut
```

View configured monitoring sessions

In the State field, `true` indicates that the port is enabled. In the Reason field, `Is UP` indicates that hardware resources are allocated.

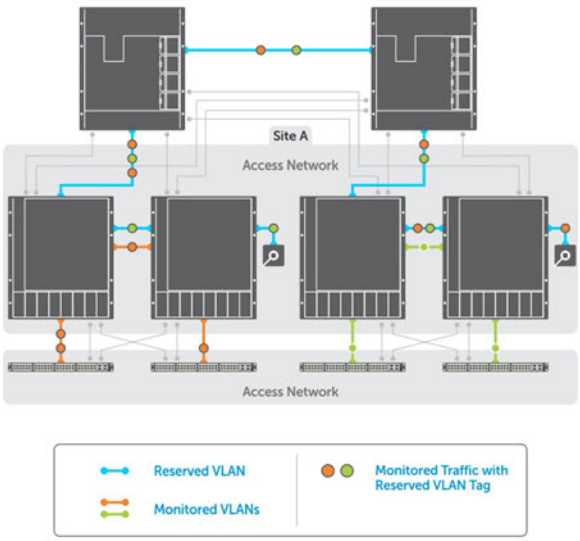
```
OS10# show monitor session all
S.Id  Source        Destination    Dir SrcIP DstIP DSCP TTL  State Reason
```

```
-------------------------------------------------------------------------
1    ethernet1/1/7  ethernet1/1/1 rx  N/A   N/A   N/A  N/A  true  Is UP
```

Remote port monitoring

Remote port monitoring monitors ingress and/or egress traffic on multiple source ports of multiple devices and forwards the monitored traffic to multiple destination ports on different remote devices. Remote port monitoring helps network administrators monitor and analyze traffic to troubleshoot network problems.

In a remote port monitoring session, monitored traffic is tagged with a VLAN ID and switched on a user-defined, non-routable L2 VLAN. The VLAN is reserved to carry only monitored traffic, which is forwarded on all egress ports of the VLAN. You must configure each intermediate switch that participates in transporting monitored traffic with the reserved L2 VLAN. Remote port monitoring supports monitoring sessions where multiple source and destination ports distribute across multiple network devices.



Session and VLAN requirements

RPM requires the following:

- Source session, such as monitored ports on different source devices.
- Reserved tagged VLAN for transporting monitored traffic configured on source, intermediate, and destination devices.
- Destination session, where destination ports connect to analyzers on destination devices.

Configure any network device with source and destination ports. Enable the network device to function in an intermediate transport session for a reserved VLAN for multiple remote port monitoring sessions. You can enable and disable individual monitoring sessions.

Consider the following when configuring a RPM session:

- A remote port monitoring session mirrors monitored traffic by prefixing the reserved VLAN tag to monitored packets to transmit using the reserved VLAN.
- The source address, destination address, and original VLAN ID of the mirrored packet are prefixed with the tagged VLAN header. Untagged source packets are tagged with the reserved VLAN ID.
- The member port of the reserved VLAN must have the MTU and IPMTU value as `MAX+4` to hold the VLAN tag parameter.
- To associate with the source session, the reserved VLAN can have up to four member ports.
- To associate with the destination session, the reserved VLAN can have multiple member ports.
- The reserved VLAN cannot have untagged ports.

Reserved L2 VLAN

- MAC address learning in the reserved VLAN is automatically disabled.
- There is no restriction on the VLAN IDs used for the reserved remote monitoring VLAN. Valid VLAN IDs are from 2 to 4093. The default VLAN ID is not supported.
- In monitored traffic, if the device has a L3 VLAN configured, packets that have the same destination MAC address as an intermediate or destination device in the path the reserved VLAN uses to transport the mirrored traffic are dropped by the device that receives the traffic .

Source session

- Configure physical ports and port-channels as sources in remote port monitoring and use them in the same source session. You can use both L2, configured with the `switchport` command, and L3 ports as source ports. Optionally, configure one or more source VLANs to configure the VLAN traffic to be monitored on source ports.
- Use the default VLAN and native VLANs as a source VLAN.
- You cannot configure the dedicated VLAN used to transport mirrored traffic as a source VLAN.

Restrictions

- When you use a source VLAN, enable flow-based monitoring using the `flow-based enable` command.
- In a source VLAN, only received (`rx`) traffic is monitored.
- In S5148F-ON, only received (`rx`) traffic is monitored.
- You cannot configure a source port-channel or source VLAN in a source session if the port-channel or VLAN has a member port configured as a destination port in a remote port monitoring session.
- You cannot use a destination port for remote port monitoring as a source port, including the session the port functions as the destination port.
- The reserved VLAN used to transport mirrored traffic must be a L2 VLAN.; L3 VLANs are not supported.

Configure remote port monitoring

Remote port monitoring requires a source interface, monitored ports on different source network devices, and a reserved tagged VLAN for transporting mirrored traffic configured on the source, intermediate, and destination devices.

1 Create a remote monitoring session in CONFIGURATION mode.

```
monitor session session-id type rpm-source
```

2 Enter the source to monitor traffic in MONITOR-SESSION mode.

```
source interface interface-range direction
```

3 Enter the destination to send the traffic to in MONITOR-SESSION mode.

```
destination remote-vlan vlan-id
```

4 Enable the monitoring interface in MONITOR-SESSION mode.

```
no shut
```

Create remote monitoring session

```
OS10(config)# monitor session 10 type rpm-source
OS10(conf-mon-rpm-source-10)#
```

Configure source and destination port, and traffic direction

```
OS10(conf-mon-rpm-source-10)# source interface vlan 10 rx
OS10(conf-mon-rpm-source-10)# destination remote-vlan 100
OS10(conf-mon-rpm-source-10)# no shut
```

View monitoring session

```
OS10(conf-mon-rpm-source-10)# do show monitor session all
S.Id  Source  Destination Dir SrcIP DstIP DSCP TTL  State Reason
```

```
-------------------------------------------------------------
1     vlan10  vlan 100   rx   N/A   N/A   N/A  N/A  true  Is UP
```

Encapsulated remote port monitoring

You can also have the monitored traffic transmitted over an L3 network to a remote analyzer. The encapsulated remote port monitoring (ERPM) session mirrors traffic from the source ports, LAGs or source VLANs and forwards the traffic using routable GRE-encapsulated packets to the destination IP address specified in the session.

Consider the following when configuring an ERPM session:

- OS10 supports only the ERPM source session. The encapsulated packets terminate at the destination IP address, the remote analyzer.
- The source IP address must be a valid local IP address for the session.
- The destination IP address must be on a remote L3 node that supports standard GRE decapsulation.
- If the destination IP address is not reachable, the session goes down.
- OS10 does not support an ERPM destination session and decapsulation of ERPM packets at the destination switch.
- You can configure a maximum of four ERPM sessions with a maximum of 128 source ports in each session. You can configure these four ERPM sessions in one of the following methods:
  - Single directional with either four ingress or four egress sessions.
  - Bidirectional with two ingress and two egress sessions.
- You can monitor a source VLAN only through flow-based monitoring. Only ingress is supported in flow-based source VLAN monitoring.
- You cannot configure an interface with ERPM traffic as a source for an ERPM session.
- You cannot monitor an RPM VLAN as a source.
- You cannot configure the same destination IP address for two sessions.
- You cannot configure an interface that serves as egress for a GRE tunnel as a source interface.
- ERPM supports only GRE-over-IPv4 tunneling.
- ERPM does not support Equal cost multi-path (ECMP).
- You can use third party devices as only tunnel-transit devices.
- OS10 does not support monitoring VLAN sub-interfaces and CPU-generated packets.

Configure encapsulated remote port monitoring

Encapsulated remote port monitoring requires valid source and destination IP addresses. Ensure that the source IP address is local and destination IP address is remote. You can also configure the time-to-live (TTL) and differentiated services code point (DSCP) values.

1 Create monitoring session in CONFIGURATION mode.

```
monitor session session-id type erpm-source
```

2 Configure source port in MONITOR-SESSION mode.

```
source interface interface-type {both | rx | tx}
```

3 Configure source and destination IP addresses, and protocol type in MONITOR-SESSION mode.

```
source-ip source ip-address destination-ip destination ip-address [gre-protocol protocol-value]
```

4 Configure TTL and DSCP values in MONITOR-SESSION mode.

```
ip {ttl ttl-number | dscp dscp-number}
```

5 Enable the monitoring interface in MONITOR-SESSION mode.

```
no shut
```

Create monitoring session

```
OS10(config)# monitor session 10 type erpm-source
OS10(conf-mon-erpm-source-10)#
```

Configure source port, source and destination IP addresses, and protocol type

```
OS10(conf-mon-erpm-source-10)# source interface ethernet 1/1/2
OS10(conf-mon-erpm-source-10)# source-ip 1.1.1.1 destination-ip 3.3.3.3 gre-protocol 35006
OS10(conf-mon-erpm-source-10)# ip ttl 16
OS10(conf-mon-erpm-source-10)# ip dscp 63
OS10(conf-mon-erpm-source-10)# no shut
```

View configured ERPM session

```
OS10(conf-mon-erpm-source-6)# do show monitor session all

S.Id  Source     Destination  Dir   Mode  Source IP  Dest IP  DSCP  TTL  Gre-Protocol State
Reason
---------------------------------------------------------------------------------------------
--
6   ethernet1/1/2 remote-ip   both  port  1.1.1.1    3.3.3.3   63    16  35006         true  Is
UP
```

View running configuration of monitor session

```
OS10# show running-configuration monitor
!
monitor session 10 type erpm-source
source-ip 1.1.1.1 destination-ip 3.3.3.3
source interface ethernet1/1/2
no shut
```

Flow-based monitoring

Flow-based monitoring conserves bandwidth by inspecting only specified traffic instead of all interface traffic. Using flow-based monitoring, you can monitor only traffic received by the source port that matches criteria in ingress access-lists (ACLs). IPv4 ACLs, IPv6 ACLs, and MAC ACLs support flow-based monitoring.

1 Enable flow-based monitoring for a monitoring session in MONITOR-SESSION mode.

```
flow-based enable
```

2 Return to CONFIGURATION mode.

```
exit
```

3 Create an access list in CONFIGURATION mode.

```
ip access-list access-list-name
```

4 Define access-list rules using `seq`, `permit`, and `deny` statements in CONFIG-ACL mode. ACL rules describe the traffic to monitor.

```
seq sequence-number {deny | permit} {source [mask] | any | host ip-address} [count [byte]]
[fragments] [threshold-in-msgs count] [capture session session-id]
```

5 Return to CONFIGURATION mode.

```
exit
```

6 Apply the flow-based monitoring ACL to the monitored source port in CONFIGURATION mode. The access list name can have a maximum of 140 characters.

```
ip access-group access-list-name {in | out}
```

Enable flow-based monitoring

```
OS10(config)# monitor session 1
OS10(conf-mon-local-1)# flow-based enable
OS10(conf-mon-local-1)# exit
OS10(config)# ip access-list ipacl1
OS10(conf-ipv4-acl)# deny ip host 1.1.1.23 any capture session 1 count
OS10(conf-ipv4-acl)# exit
OS10(config)# mac access-list mac1
OS10(conf-mac-acl)# deny any any capture session 1
OS10(conf-mac-acl)# exit
OS10(config)# interface ethernet 1/1/9
OS10(conf-if-eth1/1/9)# mac access-group mac1 in
OS10(conf-if-eth1/1/9)# end
```

```
OS10# show mac access-lists in
Ingress MAC access-list mac1
 Active on interfaces :
  ethernet1/1/9
 seq 10 deny any any capture session 1 count (0 packets)
```

Remote port monitoring on VLT

In a network, devices you configure with peer VLT nodes are considered as a single device. You can apply remote port monitoring (RPM) on the VLT devices in a network.

In a failover case, the monitored traffic reaches the packet analyzer connected to the top-of-rack (ToR) through the VLT interconnect link.

NOTE:

- In VLT devices configured with RPM, when the VLT link is down, the monitored packets might drop for some time. The time is equivalent to the VLT failover recovery time, the delay restore.
- ERPM does not work on VLT devices.

RPM on VLT scenarios

Consider a simple VLT setup where two VLT devices are connected using VLTi and a top-of-rack switch is connected to both the VLT peers using VLT LAGs in a ring topology. In this setup, the following table describes the possible scenarios when you use RPM to mirror traffic.

NOTE: Ports that connect to the VLT domain, but not part of the VLT-LAG, are called orphan ports.

Table 12. RPM on VLT scenarios

Scenario	Recommendation

Mirror an orphan port or VLT LAG or VLTi member port to a VLT LAG. The packet analyzer connects to the ToR switch.	The recommended configuration on the peer VLT device:
	1 Create a RPM VLAN
	`!`
	`interface vlan 100`
	`no shutdown`
	`remote-span`
	`!`
	2 Create an L2 ACL for the RPM VLAN - RPM session and attach it to VLTi LAG interface.
	`!`
	`mac access-list rpm`
	`seq 10 permit any any capture session 10 vlan 100`
	`!`
	`interface ethernet 1/1/1`
	`no shutdown`
	`switchport access vlan 1`
	`mac access-group rpm in`
	`!`
	3 Create a flow-based RPM session on the peer VLT device to monitor the VLTi LAG interface as the source.
	`!`
	`monitor session 10 type rpm-source destination remote-vlan 100`
	`flow-based enable`
	`source interface ethernet1/1/1 (ICL lag`

Scenario	Recommendation
	member) !
Mirror a VLAN with VLTi LAG as a member to any orphan port on the same VLT device. The packet analyzer connects to the local VLT device through the orphan port.	The recommended configuration on the VLT device: 1 Create an L2 ACL for the local session and attach it to the VLTi LAG interface. ! mac access-list local seq 10 permit any any capture session 10 ! interface ethernet 1/1/1 no shutdown switchport access vlan 1 mac access-group local in ! 2 Create a flow-based local session on the VLT device to monitor the VLTi LAG interface member (ethernet 1/1/1) as source. ! monitor session 10 type destination interface ethernet 1/1/10 flow-based enable source interface ethernet1/1/1 no shut !
Mirror a VLAN with a VLTi LAG as the member to the VLT LAG on the same VLT device. The packet analyzer connects to the ToR switch.	—
Mirror a VLT LAG of the ToR, or any port in the ToR to any orphan port in the VLT device. Configure VLT nodes as intermediate devices. The packet analyzer connects to the ToR switch.	—
Mirror a VLT LAG to any orphan port on the same VLT device. The packet analyzer connects to the local VLT device through the orphan port.	If the packet analyzer directly connects to the VLT peer where the source session is configured, use local port monitoring instead of RPM.
Mirror an orphan port in the primary VLT device to any orphan port on a secondary VLT device through the VLTi. The packet analyzer connects to the secondary VLT device through the orphan port. In this case, the mirroring packets duplicate.	—
Mirror a VLT LAG of the primary VLT device to any orphan port on a secondary VLT device through the VLTi. The packet analyzer connects to the secondary VLT device through the orphan port.	—
Mirror a member port of the VLTi LAG or VLT LAG to any orphan port in the same device. The packet analyzer connects to the local VLT device through the orphan port.	If the packet analyzer is directly connected to the VLT peer in which the source session is configured, use local port monitoring instead of RPM.
Mirror a member port of VLTi LAG to the VLT LAG on the same VLT device. The packet analyzer connects to the ToR switch.	—
Mirror a VLT LAG or VLT member port as part of the source VLAN and destination VLAN. The packet analyzer connects to the ToR switch.	—

Port monitoring commands

description

Configures a description for the port monitoring session. The monitoring session can be: local, RPM, or ERPM.

Syntax `description string`

Parameters *string* — Enter a description of the monitoring session. A maximum of 255 characters.

Default Not configured

Command Mode MONITOR-SESSION

Usage Information The `no` version of this command removes the description text.

Example

```
OS10(conf-mon-local-1)# description remote
OS10(conf-mon-rpm-source-5)# description "RPM Sesssion"
OS10(conf-mon-erpm-source-10)# description "ERPM Session"
```

Supported Releases 10.2.0E or later

destination

Sets the destination where monitored traffic is sent to. The monitoring session can be local or RPM.

Syntax `destination {interface interface-type | remote-vlan vlan-id}`

Parameters *interface-type* — Enter the interface type for a local monitoring session.

- `ethernet node/slot/port[:subport]` — Enter the Ethernet interface information as the destination.
- `port-channel id-number` — Enter a port-channel number as the destination, from 1 to 128.
- `vlan vlan-id` —Enter a VLAN ID as the destination, from 1 to 4093.

`remote-vlan vlan-id`—Enter a remote VLAN ID as the destination for the RPM monitoring session, from 1 to 4093.

Default Not configured

Command Mode MONITOR-SESSION

Usage Information The `no` version of this command resets the value to the default.

Example

```
OS10(conf-mon-local-10)# destination interface port-channel 10
OS10(conf-mon-rpm-source-3)# destination remote-vlan 20
```

Supported Releases 10.2.0E or later

flow-based

Enables flow-based monitoring. The monitoring session can be: local, RPM, or ERPM.

Syntax `flow-based enable`

Parameters None

Default Disabled

Command Mode MONITOR-SESSION

Usage Information The `no` version of this command disables the flow-based monitoring.

Example

```
OS10(conf-mon-local-1)# flow-based enable
OS10(conf-mon-rpm-source-2)# flow-based enable
OS10(conf-mon-erpm-source-3)# flow-based enable
```

Supported Releases 10.2.0E or later

ip

Configures the IP time-to-live (TTL) value and the differentiated services code point (DSCP) value for the ERPM traffic.

Syntax `ip {ttl ttl-number | dscp dscp-number}`

Parameters

- *ttl-number* — Enter the TTL value, from 1 to 255.
- *dscp-number* — Enter the DSCP value, from 0 to 63.

Default

- TTL: 255
- DSCP: 0

Command Mode MONITOR-SESSION (ERPM)

Usage Information The `no` version of this command removes the configured TTL and DSCP values.

Example

```
OS10(conf-mon-erpm-source-10)#  ip ttl 16
OS10(conf-mon-erpm-source-10)#  ip DSCP 63
```

Supported Releases 10.4.0E(R1) or later

monitor session

Creates a session for monitoring traffic with port monitoring.

Syntax `monitor session session-id type [local | rpm-source | erpm-source]`

Parameters

- *session-id* — Enter a monitor session ID, from 1 to 18.
- `local` — (Optional) Enter a local monitoring session.
- `rpm-source` — (Optional) Enter a remote monitoring session.
- `erpm-source` — (Optional) Enter an encapsulated remote monitoring session.

Default local

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the monitor session.

Example
```
OS10(config)# monitor session 1
OS10(conf-mon-local-1)#
```

Example (RPM)
```
OS10(config)# monitor session 5 type rpm-source
OS10(conf-mon-rpm-source-5)#
```

Example (ERPM)
```
OS10(config)# monitor session 10 type erpm-source
OS10(conf-mon-erpm-source-10)#
```

Supported Releases 10.2.0E or later

show monitor session

Displays information about a monitoring session.

Syntax `show monitor session {session-id | all}`

Parameters

- *session-id* — Enter the session ID number, from 1 to 18.
- `all` — View all monitoring sessions.

Default All

Command Mode EXEC

Usage Information In the State field, `true` indicates that the port is enabled. In the Reason field, `Is UP` indicates that hardware resources are al

Example (specific session)
```
OS10# show monitor session 1
S.Id Source      Destination Dir   Mode  Source IP   Dest IP     DSCP  TTL  Gre-Protoco
-------------------------------------------------------------------------------------
1   ethernet1/1/1 remote-ip  both  port  11.11.11.1  11.11.11.11    0  255  35006
```

Example (all sessions)
```
OS10# show monitor session all
S.Id  Source        Destination  Dir  Mode  Source IP    Dest IP     DSCP  TTL  Gre-Pro
-------------------------------------------------------------------------------------
1    ethernet1/1/1  remote-ip    both  port  11.11.11.1  11.11.11.11  0    255  35006
9    ethernet1/1/9               both  port  N/A         N/A               N/A
7    ethernet1/1/9    vlan40     both  port  N/A         N/A               N/A
4    ethernet1/1/1               both  port  N/A         N/A          0    255  35006
Destination is not resolved
6    ethernet1/1/2  remote-ip    both  port  11.11.11.1  2.2.2.1      0    255  35006
session does not exist
```

Supported Releases 10.2.0E or later

shut

Disables the monitoring session. The monitoring session can be: local, RPM, or ERPM.

Syntax `shut`

Parameters None

Default Disabled

Command Mode MONITOR-SESSION

Usage Information The `no` version of this command enables the monitoring session.

Example

```
OS10(config)# monitor session 1
OS10(conf-mon-local-1)# no shut
```

```
OS10(config)# monitor session 5 type rpm-source
OS10(conf-mon-rpm-source-5)# no shut
```

```
OS10(config)# monitor session 10 type erpm-source
OS10(conf-mon-erpm-source-10)# no shut
```

Supported Releases 10.2.0E or later

source

Configures a source for port monitoring. The monitoring session can be: local, RPM, or ERPM.

Syntax `source interface interface-type {both | rx | tx}`

Parameters

- *interface-type* — Enter the interface type:
  - `ethernet` *node/slot/port[:subport]* — Enter the Ethernet interface information as the monitored source.
  - `port-channel` *id-number* — Enter the port-channel interface number as the monitored source, from 1 to 128. This option is not supported in S5148F–ON .
  - `vlan` *vlan-id* —Enter the VLAN identifier as the monitored source, from 1 to 4093.
- `both` — Monitor both receiving and transmitting packets. This option is not supported in the S5148F–ON platform, and on VLAN interfaces for other platforms.
- `rx` — Monitor only received packets.
- `tx` — Monitor only transmitted packets. This option is not supported in the S5148F–ON platform, and on VLAN interfaces for other platforms.

Default Not configured

Command Mode MONITOR-SESSION

Usage Information None

Example

```
OS10(config)# monitor session 1
OS10(conf-mon-local-1)# source interface ethernet 1/1/7 rx
```

```
OS10(config)# monitor session 5 type rpm-source
OS10(conf-mon-rpm-source-5)# source interface ethernet 1/1/10 rx
```

```
OS10(config)# monitor session 10 type erpm-source
OS10(conf-mon-erpm-source-10)# source interface ethernet 1/1/5 rx
```

Supported Releases 10.2.0E or later

source-ip

Configures the source, destination, and protocol type of the monitored port for an ERPM monitoring session.

Syntax `source-ip` *source ip-address* `destination-ip` *destination ip-address* `[gre-protocol` *protocol-value*`]`

Parameters

- *source ip-address* — Enter the source IP address.
- *destination ip-address* — Enter the destination IP address.
- *protocol-value* — Enter the GRE protocol value, from 1 to 65535, default: 35006.

Default Not configured

Command Mode MONITOR-SESSION

Usage Information None

Example

```
OS10(config)# monitor session 10
OS10(conf-mon-erpm-source-10)# source-ip 10.16.132.181 destination-ip
172.16.10.11 gre-protocol 35006
```

Supported Releases 10.4.0E(R1) or later

5

Layer 3

Bidirectional forwarding detection (BFD)	Provides rapid failure detection in links with adjacent routers (see BFD commands).
Border Gateway Protocol (BGP)	Provides an external gateway protocol that transmits inter-domain routing information within and between autonomous systems (see BGP Commands).
Equal Cost Multi-Path (ECMP)	Provides next-hop packet forwarding to a single destination over multiple best paths (see ECMP Commands).
IPv4 Routing	Provides forwarding of packets to a destination IP address, based on a routing table. This routing table defines how packets are routed — dynamically, broadcasted directly to, using proxy ARP, as well as what type of information is included with the packets (see IPv4 Routing Commands).
IPv6 Routing	Provides routing for the IPv6 address space, stateless auto-configuration, header format simplifications, and improved support for options and extensions (see IPv6 Routing Commands).
Open Shortest Path First (OSPF)	Provides a link-state routing protocol that communicates with all other devices in the same autonomous system area using link-state advertisements (LSAs). OS10 supports up to 10,000 OSPF routes for OSPFv2 to designate up to 8,000 routes as external, and up to 2,000 as inter/intra area routes (see OSPF Commands).
Virtual Router Redundancy Protocol (VRRP)	Provides a mechanism to eliminate a single point of failure in a statically routed network (see VRRP Commands).
Virtual Routing and Forwarding (VRF)	Provides a mechanism to partition a physical router into multiple virtual routers (see VRF Commands).

Virtual routing and forwarding

VRF partitions a physical router into multiple virtual routers (VRs). The control and data plane are isolated in each VR; traffic does not flow across VRs. VRF allows multiple instances of routing tables to co-exist within the same router at the same time.

OS10 supports a management VRF instance, a default VRF instance, and a maximum of 512 non-default VRF instances. Use the default and non-default VRF instances to configure routing.

You can move the management interface from the default to management VRF instance. You need not create the management VRF instance as it already exists in the system by default.

By default, OS10 initially assigns all physical interfaces and all logical interfaces to the default VRF instance.

Configure management VRF

You can assign only management interfaces to the management VRF instance.

Before you assign the management interface to the management VRF instance, remove all the configured settings, including the IP address, on the management interface.

1. Enter the `ip vrf management` command in CONFIGURATION mode. Use Non-Transaction-Based Configuration mode only. Do not use Transaction-Based mode.
2. Add the management interface using the `interface management` command in VRF CONFIGURATION mode.

Configure management VRF

```
OS10(config)# ip vrf management
OS10(conf-vrf)# interface management
```

You can enable various services in both management or default VRF instances. The services supported in the management and default VRF instances are:

Table 13. Services supported

Application	Management VRF	Default VRF	Non-default VRF
BGP	No	Yes	Yes
COPP ACL	Yes	Yes	No
DHCP client	Yes	Yes	Yes
DHCP relay	No	Yes	Yes
DHCP server	No	Yes	No
DNS client	Yes	Yes	Yes
FTP client	Yes	Yes	Yes
HTTP client	Yes	Yes	Yes
ICMP / Ping	Yes	Yes	Yes
NTP client	Yes	Yes	Yes
NTP server	Yes	Yes	Yes
OSPFV2 /OSPFV3	No	Yes	Yes
RADIUS server	Yes	Yes	Yes
SCP client	Yes	Yes	Yes
sFlow	Yes	Yes	Yes
SFTP	Yes	Yes	Yes
SNMP traps	Yes	Yes	No
SSH server`	Yes	Yes	Yes
Syslog	Yes	Yes	Yes
TACACS+ server	Yes	Yes	Yes
Telnet server	Yes	Yes	Yes
TFTP client	Yes	Yes	Yes
Traceroute	Yes	Yes	Yes

Application	Management VRF	Default VRF	Non-default VRF
VLT backup link	Yes	Yes	No
VRRP	Yes	Yes	Yes

Configure a static route for a management VRF instance

- Configure a static route that directs traffic to the management interface.
  CONFIGURATION

  `management route` *`ip-address mask`* `managementethernet` or `management route` *`ipv6-address prefix-length`* `managementethernet`

  You can also configure the management route to direct traffic to a physical interface. For example: `management route 10.1.1.5/24 ethernet 1/1/4` or `management route 2::/64 ethernet 1/1/2`.
- Configure a static entry in the IPv6 neighbor discovery.
  CONFIGURATION

  ```
  ipv6 neighbor vrf management 1::1 ethernet 1/1/2 xx:xx:xx:xx:xx:xx
  ```

Configure non-default VRF instances

In addition to a management VRF instance and default VRF, OS10 also supports non-default VRF instances. You can create a maximum of 128 non-default VRF instances.

While you can assign management interfaces only to the management VRF instance, you can assign any physical or logical interface – VLAN, port channel or loopback, to a non-default VRF instance.

When you create a new non-default VRF instance, OS10 does not assign any interface to it. You can assign the new VRF instance to any of the existing physical or logical interfaces, provided they are not already assigned to another non-default VRF.

NOTE: When you create a new logical interface, OS10 assigns it automatically to the default VRF instance. In addition, OS10 initially assigns all physical Layer 3 interfaces to the default VRF instance.

You can reassign any interface assigned to a non-default VRF instance back to the default VRF instance.

- Create a non-default VRF instance by specifying a name and enter VRF configuration mode.
  CONFIGURATION

  ```
  ip vrf vrf-name
  ```

Assign an interface to a non-default VRF instance

After creating a non-default VRF instance you can associate an interface to the VRF instance that you created.

To assign an interface to a non-default VRF, perform the following steps:

1. Enter the interface that you want to assign to a non-default VRF instance.
   CONFIGURATION

   ```
   interface ethernet 1/1/1
   ```
2. Remove the interface from L2 switching.
   INTERFACE

```
no switchport
```

3 Assign the interface to a non-default VRF.
INTERFACE CONFIGURATION

```
ip vrf forwarding vrf-test
```

Before assigning an interface to a VRF instance, ensure that no IP address is configured on the interface.

4 Assign an IPv4 address to the interface.
INTERFACE CONFIGURATION

```
ip address 10.1.1.1/24
```

5 Assign an IPv6 address to the interface.
INTERFACE CONFIGURATION

```
ipv6 address 1::1/64
```

You can also auto configure an IPv6 address using the `ipv6 address autoconfig` command.

NOTE: Before configuring any routing protocol in a VRF instance, you need to first assign an IP address to at least one of the interfaces assigned to the VRF instance on which you want to configure routing protocols.

Assigning a loopback interface to a non-default VRF instance

After creating a non-default VRF instance you can associate a loopback interface to the VRF instance that you created.
To assign a loopback interface to a non-default VRF, perform the following steps:

1 Enter the loopback interface that you want to assign to a non-default VRF instance.
CONFIGURATION

```
interface loopback 5
```

2 Assign the interface to a non-default VRF.
INTERFACE CONFIGURATION

```
ip vrf forwarding vrf-test
```

Before assigning a n interface to a VRF instance, ensure that no IP address is configured on the interface.

3 Assign an IPv4 address to the interface.
INTERFACE CONFIGURATION

```
ip address 10.1.1.1/24
```

4 Assign an IPv6 address to the interface.
INTERFACE CONFIGURATION

```
ipv6 address 1::1/64
```

You can also auto configure an IPv6 address using the `ipv6 address autoconfig` command.

Assign an interface back to the default VRF instance

Table 14. Configurations to be removed

CONFIGURATION	MODE	COMMAND
IP address — In interface configuration mode, undo the IP address configuration.	INTERFACE CONFIGURATION	`OS10(conf-if-eth1/1/10:1)#no ip address` *ipv4-address* or `no ipv6 address` *ipv6-address*
Port — In interface configuration mode, remove the interface association corresponding to the VRF instance that you want to delete.	INTERFACE CONFIGURATION	`OS10(conf-if-eth1/1/10:1)#no ip vrf forwarding`

To assign an interface back to the default VRF, perform the following steps:

1 Enter the interface that you want to assign back to the default VRF instance.
CONFIGURATION

```
interface ethernet 1/1/1
```

2 Remove the IPv4 address associated with the interface.
INTERFACE CONFIGURATION

```
no ip address
```

3 Remove the IPv6 address associated with the interface.
INTERFACE CONFIGURATION

```
no ipv6 address
```

4 Assign the interface back to the default VRF instance.
INTERFACE CONFIGURATION

```
no ip vrf forwarding
```

Assigning the management interface back to the default VRF instance

To assign the management interface back to the default VRF, perform the following steps:

1 Enter the management VRF instance.
CONFIGURATION

```
ip vrf management
```

2 Assign the management interface back to the default VRF instance.
CONFIGURATION VRF

```
no interface management
```

Deleting a non-default VRF instance

Before deleting a non-default VRF instance, ensure all the dependencies and associations corresponding to that VRF instance are first removed or disabled. Following table shows the dependencies that you have to remove before deleting a non-default VRF instance: After removing all dependences, you can delete the non-default VRF instances that you create.

- Delete a non-default VRF instance using the following command:
  CONFIGURATION

  ```
  no ip vrf vrf-name
  ```

  NOTE: You cannot delete the default VRF instance.

Configure a static route for a non-default VRF instance

- Configure a static route in a non-default VRF instance. Static routes contain IP addresses of the next-hop neighbors that are reachable through the non-default VRF. These IP addresses could also belong to the interfaces that are part of the non-default VRF instance.
  CONFIGURATION

  `ip route vrf vrf-name ip-address mask next-hop-ip-address` or `ipv6 route vrf vrf-name ipv6-address prefix-length next-hp[=ipv6-address`

  For example: `ip route vrf red 10.1.1.0/24 20.1.1.6 or ipv6 route vrf red 2::/64 3::1`
- Configure the route to direct traffic to a front-panel port in case of a non-default VRF instance.
  CONFIGURATION

  `ip route ip-address-mask ethernet interface-type or ipv6 route ipv6-address-mask ethernet interface-type`

  For example: `ip route 10.1.1.0/24 ethernet 1/1/1` or `ipv6 route 2::/64 ethernet 1/1/1`. Where ethernet 1/1/1 is part of the non-default VRF.

Configuring static entry in IPv6 neighbor

- Configure a static entry in the IPv6 neighbor discovery.
  CONFIGURATION

  ```
  ipv6 neighbor vrf vrf-test 1::1 ethernet 1/1/1 xx:xx:xx:xx:xx:xx
  ```

VRF configuration

The following configuration illustrates a typical VRF setup:

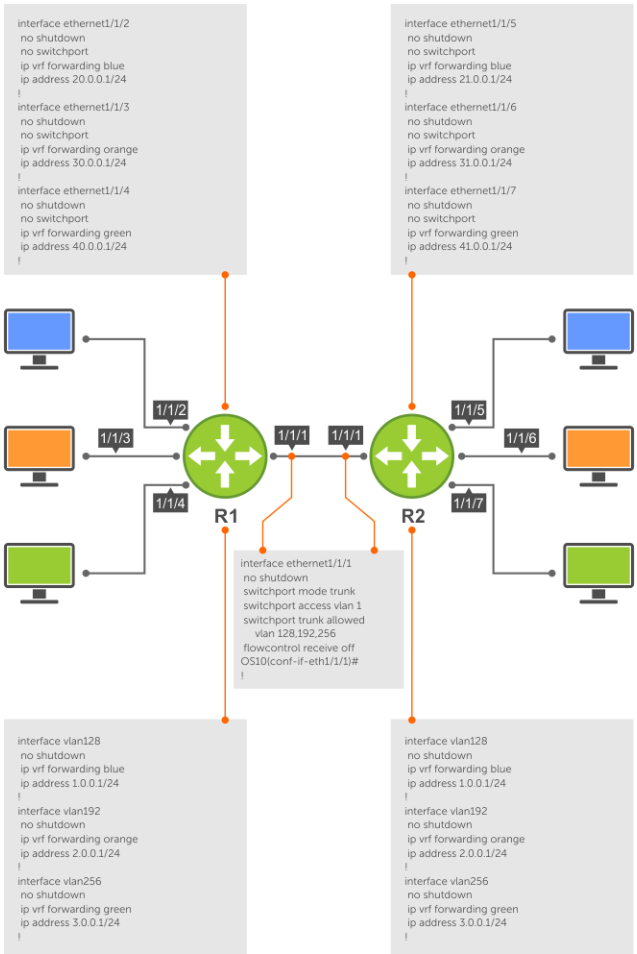


Figure 3. Setup VRF Interfaces

Router 1

```
ip vrf blue
!
ip vrf orange
!
ip vrf green
!
interface ethernet 1/1/1
    no ip address
    no switchport
    no shutdown
!
interface ethernet1/1/2
 no shutdown
 no switchport
 ip vrf forwarding blue
 ip address 20.0.0.1/24
!
interface ethernet1/1/3
 no shutdown
 no switchport
 ip vrf forwarding orange
 ip address 30.0.0.1/24
!
interface ethernet1/1/4
```

```
 no shutdown
 no switchport
 ip vrf forwarding green
 ip address 40.0.0.1/24
!
interface vlan128
 mode L3
 no shutdown
 ip vrf forwarding blue
 ip address 1.0.0.1/24
!
interface vlan192
 mode L3
 no shutdown
 ip vrf forwarding orange
 ip address 2.0.0.1/24
!
!
interface vlan256
 mode L3
 no shutdown
 ip vrf forwarding green
 ip address 3.0.0.1/24
!
ip route vrf green 30.0.0.0/24 3.0.0.1
```

Router 2

```
ip vrf blue
!
ip vrf orange
!
ip vrf green
!
interface ethernet 1/1/1
    no ip address
    no switchport
    no shutdown
!
interface ethernet1/1/5
 no shutdown
 no switchport
 ip vrf forwarding blue
 ip address 21.0.0.1/24
!
interface ethernet1/1/6
 no shutdown
 no switchport
 ip vrf forwarding orange
 ip address 31.0.0.1/24
!
interface ethernet1/1/7
 no shutdown
 no switchport
 ip vrf forwarding green
 ip address 41.0.0.1/24
!
interface vlan128
 mode L3
 no shutdown
 ip vrf forwarding blue
 ip address 1.0.0.1/24
!
interface vlan192
 mode L3
 no shutdown
 ip vrf forwarding orange
 ip address 2.0.0.1/24
!
```

```
interface vlan256
 mode L3
 no shutdown
 ip vrf forwarding green
 ip address 3.0.0.1/24
!
ip route vrf green 31.0.0.0/24 3.0.0.1
```

Router 1 show command output

```
OS10# show ip vrf
VRF-Name                          Interfaces
blue                              Eth1/1/2
                                  Vlan128

default                           Mgmt1/1/1
                                  Vlan1,24-25,200

green                             Eth1/1/4
                                  Vlan256

orange                            Eth1/1/3
                                  Vlan192

OS10# show ip route vrf blue
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, * - candidate default,
       + - summary route, > - non-active route
Gateway of last resort is not set
  Destination         Gateway                          Dist/Metric  Last Change
--------------------------------------------------------------------------------
  C     20.0.0.0/24   via 20.0.0.1    ethernet1/1/2   0/0         01:46:41

OS10# show ip route vrf orange
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, * - candidate default,
       + - summary route, > - non-active route
Gateway of last resort is not set
  Destination         Gateway                           Dist/Metric  Last Change
--------------------------------------------------------------------------------
  C     30.0.0.0/24   via 30.0.0.1    ethernet1/1/3    0/0          01:55:00

OS10# show ip route vrf green
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, * - candidate default,
       + - summary route, > - non-active route
Gateway of last resort is not set
  Destination         Gateway                            Dist/Metric  Last Change
--------------------------------------------------------------------------------
  C     40.0.0.0/24   via 40.0.0.1    ethernet1/1/4     0/0           02:01:15
================================================================================
```

Router 2 show command output

```
OS10# show ip vrf
VRF-Name                          Interfaces
blue                              Eth1/1/5
```

```
                                  Vlan128

default                           Mgmt1/1/1
                                  Vlan1,24-25,200

green                             Eth1/1/7
                                  Vlan256

orange                            Eth1/1/6
                                  Vlan192

OS10# show ip route vrf blue
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, * - candidate default,
       + - summary route, > - non-active route
Gateway of last resort is not set
  Destination         Gateway                          Dist/Metric  Last Change
-------------------------------------------------------------------------------
  C     21.0.0.0/24   via 21.0.0.1    ethernet1/1/5    0/0          02:05:00

OS10# show ip route vrf orange
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, * - candidate default,
       + - summary route, > - non-active route
Gateway of last resort is not set
  Destination         Gateway                           Dist/Metric  Last Change
-------------------------------------------------------------------------------
  C     31.0.0.0/24   via 31.0.0.1    ethernet1/1/6     0/0          02:09:19

OS10# show ip route vrf green
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, * - candidate default,
       + - summary route, > - non-active route
Gateway of last resort is not set
  Destination          Gateway                          Dist/Metric  Last Change
-------------------------------------------------------------------------------
  C     41.0.0.0/24    via 41.0.0.1   ethernet1/1/7     0/0          02:45:16
===============================================================================
```

View VRF instance information

To display information about a VRF configuration, use the `show ip vrf` command. To display information on all VRF instances including the default VRF 0, do not enter a value for *vrf-name*.

- Display the VRF instance interfaces.
  EXEC

  `show ip vrf [vrf-name]`

Static route leaking

Route leaking enables routes that are configured in a default or non-default VRF instance to be made available to another VRF instance. You can leak routes from a source VRF instance to a destination VRF instance.

The routes need to be leaked in both source as well as destination VRFs in order to achieve end-to-end traffic flow.

If there are any connected routes in the same subnet as statically leaked routes. then the connected routes take precedence.

In static route leaking, DHCP functionality does not work for overlapping subnets. For example, if two interfaces on different VRFs are on the same subnet and are configured with the same DHCP server, then only one of those interface get an IP address. This is because the client requests from these interfaces have the same MAC and subnet addresses. The server does not have any unique parameter to differentiate that the request is from two different clients.

Limitations

- In VLT scenarios, the resolved ARP entry for the leaked route is not synced between the VLT peers. The ARP entry resolved in the source VRF is programmed into the leaked VRF when the leaked route configuration is active.
- During downgrade from 10.4.2, the leaked route configuration is restored. However, the routes remain inactive in the destination VRF instance.
- During downgrade from 10.4.2, the **update-source-if** command is not restored.

Configuring static route leaking

To configure static route leaking:

1 Enter the interface in the source VRF instance that contains the static routes that you want to leak.

```
interface interface-name
```

CONFIGURATION Mode

2 In INTERFACE CONFIGURATION Mode, assign the interface to the source VRF instance.

```
ip vrf forwarding vrf1
```

INTERFACE CONFIGURATION Mode

3 Assign an IP address to the interface.

```
ip address ip-address
```

VRF CONFIGURATION Mode

4 Enter the interface of the VRF instance to which you want to leak the static routes.

```
interface interface-name
```

CONFIGURATION Mode

5 In INTERFACE CONFIGURATION Mode, assign the interface to the destination VRF instance.

```
ip vrf forwarding vrf2
```

INTERFACE CONFIGURATION Mode

6 Configure the static route that you want to leak on the destination VRF instance.

```
ip route vrf dest-vrf-name route nexthop-interface
```

7 Configure the static route that you have configured earlier in the source VRF instance to be available in the destination VRF instance also.

```
ip route vrf src-vrf-name route nexthop-interface
```

```
OS10(config)#interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ip vrf forwarding VRF1
OS10(conf-if-eth1/1/1)# ip address 120.0.0.1/24
OS10(config)#interface ethernet 1/1/2
OS10(conf-if-eth1/1/1)# ip vrf forwarding VRF2
OS10(conf-if-eth1/1/1)# ip address 140.0.0.1/24
OS10(config)#ip route vrf VRF1 140.0.0.0/24 interface ethernet 1/1/2
OS10(config)#ip route vrf VRF2 120.0.0.0/24 interface ethernet 1/1/1

The following example shows the show output:
OS10(config)# do show ip route vrf VRF1
Codes: C - connected
 S - static
 B - BGP, IN - internal BGP, EX - external BGP
 O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
 N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
 E2 - OSPF external type 2, * - candidate default,
 + - summary route, > - non-active route
Gateway of last resort is not set
 Destination Gateway Dist/Metric Last Change
-------------------------------------------------------------------------------------------------
----
 C 120.0.0.0/24 via 120.0.0.1 ethernet1/1/1 0/0 00:00:57
 S 140.0.0.0/24 Direct,VRF2 ethernet1/1/2 1/0 00:00:04

OS10(config)# do show ip route vrf VRF2
Codes: C - connected
 S - static
 B - BGP, IN - internal BGP, EX - external BGP
 O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
 N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
 E2 - OSPF external type 2, * - candidate default,
 + - summary route, > - non-active route
Gateway of last resort is not set
 Destination Gateway Dist/Metric Last Change
-------------------------------------------------------------------------------------------------
----
 S 120.0.0.0/24 Direct,VRF1 ethernet1/1/1 1/0 00:00:05
 C 140.0.0.0/24 via 140.0.0.1 ethernet1/1/2 0/0 00:01:54
```

Configuring source IP address for a leaked route

If the source IP is not mentioned explicitly for any self originating packet (for example, ping or traceroute) to the leaked route destined through the parent VRF, the system chooses a source based on its source selection algorithm.

NOTE: For end-to-end traffic to flow, you must specify the source for self originating packets and leak the same into the destination VRF.

To mitigate this issue and have control over the source IP address for leaked routes, you can create a loopback interface and associate it with the leaked VRF.

To explicitly mention the source interface for the leaked VRF:

Enter the following command:

```
update-source-if
```

VRF CONFIGURATION Mode

After you configure the source IP address in a leaked VRF, if ping is initiated without -I option, then the source IP address will be that of loopback interface.

Route leaking using route targets

You can leak routes in one VRF instance to another using route targets.

NOTE: You can leak routes using route targets only on the default and non-default VRF instance. You cannot leak routes using route targets on the management VRF instance.

To leak routes in one VRF instance using route targets:

1 Enter the VRF from which you want to leak routes using route targets.
CONFIGURATION

```
ip vrf source-vrf-name
```

2 Export the route that belongs to one VRF instance.
VRF CONFIGURATION

```
ip route-export 1::1
```

3 Enter the non-default VRF instance to which you want to leak the route.
CONFIGURATION

```
ip vrf destination-vrf-name
```

```
ip route-import 1::1
```

The routes that you exported from the source VRF instance are now available in the destination VRF instance.

NOTE: When you configure static route leaking, only a single ECMP route- the best ECMP route - leaks to another VRF instance. The other ECMP static routes are not leaked to the other VRF instances.

Route leaking using route maps

You can leak routes in one VRF instance to another VRF instance using route maps.
To leak routes in one VRF instance using route maps:

1 Enter the VRF from which you want to leak routes using route targets.
CONFIGURATION

```
ip vrf source-vrf-name
```

```
ip vrf VRF-A
```

2 Configure the IP prefix.
CONFIGURATION

```
ip prefix-list prefix-list-name {permit | deny} ip-address
```

```
ip prefix-list abc permit 20.0.0.0/24
or
ip prefix-list abc deny 20.0.0.0/24
```

3 Configure the route-map.
CONFIGURATION

```
route-map route-map-name

route-map xyz
```

4 Associate the prefix list to the route-map.

CONFIGURATION

```
route-map route-map-name {permit | deny} rule

match ip address prefix-list prefix-list-name

route-map xyz permit 10
match ip address prefix-list abc
or
route-map xyz deny 10
match ip address prefix-list abc
```

5 Export the routes from a VRF instance using route maps.

VRF CONFIGURATION

```
ip route-export route-target route-map-name route-map-name

ip route-export 1:1 route-map xyz
```

6 Import these leaked routes into another VRF instance.

VRF CONFIGURATION

```
ip route-import route-target

ip route-import 1:1
```

Example Configuration for IPv4 static route leaking

```
OS10(config)#interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ip vrf forwarding VRF1
OS10(conf-if-eth1/1/1)# ip address 120.0.0.1/24

OS10(config)#interface ethernet 1/1/2
OS10(conf-if-eth1/1/2)# ip vrf forwarding VRF2
OS10(conf-if-eth1/1/2)# ip address 140.0.0.1/24
OS10(config)#ip route vrf VRF1 160.0.0.0/24 120.0.0.2
OS10(config)# ip vrf VRF1
OS10(conf-vrf)# ip route-export 1:1
OS10(config)# ip vrf VRF2
OS10(conf-vrf)# ip route-import 1:1
```

VRF commands

interface management

Adds a management interface to the management VRF instance.

Syntax	`interface management`
Parameters	None
Default	Not configured
Command Mode	VRF CONFIGURATION
Usage Information	The `no` version of this command removes the management interface from the management VRF instance.

Example

```
OS10(config)# ip vrf management
OS10(conf-vrf)# interface management
```

Supported Releases 10.4.0E(R1) or later

ip domain-list vrf

Configures a domain list for the management VRF instance or any non-default VRF instance that you create.

Syntax `ip domain-list vrf {management | vrf-name} domain-names`

Parameters

- `management`—Enter the keyword `management` to configure a domain list for the management VRF instance.
- *`vrf-name`*—Enter the name of the non-default VRF instance to configure a domain list for that non-default VRF instance.
- *`domain-names`*—Enter the list of domain names.

Default Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the domain list configuration from the management or the non-default VRF instance.

Example

```
OS10(config)# ip domain-list vrf management dns1 dell.com
or
OS10(config)# ip domain-list vrf blue dnsl dell.com
```

Supported Releases 10.4.0E(R1) or later

ip domain-name vrf

Configures a domain name for the management VRF instance or any non-default VRF instance that you create.

Syntax `ip domain—name vrf {management | vrf-name} domain-name`

Parameters

- `management`—Enter the keyword `management` to configure a domain name for the management VRF instance.
- *`vrf-name`*—Enter the name of the non-default VRF instance to configure a domain name for that VRF instance.
- *`domain-name`*—Enter the domain name.

Default Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the domain name from the management or non-default VRF instance.

Example

```
OS10(config)# ip domain-name vrf management dell.com
or
OS10(config)# ip domain-name vrf blue dell.com
```

Supported Releases 10.4.0E(R1) or later

ip vrf

Create a non-default VRF instance.

Syntax `ip vrf vrf-name`

Parameters

- *vrf-name*—Enter the name of the non-default VRF that you want to create. Enter a VRF name that is not greater than 32 characters in length.

Default Not configured

Command Mode CONFIGURATION

Usage Information Enter the `ip vrf vrf-name` command only in non-transaction-based configuration mode. Do not use transaction-based mode. You can create up to a maximum of 128 non-default VRF instances. The `no ip vrf vrf-name` command removes the non-default VRF instance that you specify.

Example

```
OS10(config)# ip vrf vrf-test
OS10(conf-vrf-test)#
```

Supported Releases 10.4.1.0 or later

ip ftp vrf

Configures an FTP client for the management or non-default VRF instance.

Syntax `ip ftp vrf {management | vrf vrf-name}`

Parameters

- `management` — Enter the keyword to configure an FTP client on the management VRF instance.
- `vrf vrf-name` — Enter the keyword then the name of the VRF to configure an FTP client on that non-default VRF instance.

Default Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the management VRF instance configuration from the FTP client.

Example

```
OS10(config)# ip ftp vrf management
OS10(config)# ip ftp vrf vrf-blue
```

Supported Releases 10.4.0E(R1) or later

ip host vrf

Configures a host name for the management VRF instance or a non-default VRF instance and maps the host name to an IPv4 or IPv6 address.

Syntax `ip host vrf {management | vrf-name} hostname {IP-address | Ipv6-address}`

Parameters

- `management`—Enter the keyword `management` to configure a host name for the management VRF instance.
- *vrf-name*—Enter the name of the non-default VRF instance to configure a host name for that VRF instance.

- *hostname*—Enter the host name.
- *IP-address | Ipv6-address*—Enter the host IPv4 or IPv6 address.

Default Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the host name from the management or non-default VRF instance.

Example

```
OS10(config)# ip host vrf management dell 10.1.1.1
or
OS10(config)# ip host vrf blue dell 10.1.1.1
```

Supported Releases 10.4.0E(R1) or later

ip http vrf

Configures an HTTP client for the management or non-default VRF instance.

Syntax `ip http vrf {management | vrf vrf-name}`

Parameters

- `management` — Enter the keyword to configure an HTTP client for the management VRF instance.
- `vrf vrf-name` — Enter the keyword then the name of the VRF to configure an HTTP client for that non-default VRF instance.

Default Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the management VRF instance configuration from the HTTP client.

Example

```
OS10(config)# ip http vrf management
OS10(config)# ip http vrf vrf-blue
```

Supported Releases 10.4.0E(R1) or later

ip name-server vrf

Configures a DNS name server for the management VRF instance or a non-default VRF instance.

Syntax `ip name-server vrf {management | vrf-name}`

Parameters

- `management`—Enter the keyword `management` to configure a DNS name server for the management VRF instance.
- *vrf-name*—Enter the name of the non-default VRF instance to configure a DNS name server for that VRF instance.

Default Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the name server from the management or non-default VRF instance.

Example

```
OS10(config)# ip name-server vrf management
or
OS10(config)# ip name-server vrf blue
```

Supported Releases 10.4.0E(R1) or later

ip route-import

Imports an IPv4 static route into a VRF instance from another VRF instance.

Syntax

```
[no] ip route-import route-target
```

Parameters

- *route-target* — Enter the `route-target` of the non-default VRF instance, from 1 to 65535.

Default Not configured

Command Mode VRF CONFIG

Usage Information You can import routes corresponding only to a non-default or a default VRF instance. You cannot import routes that belong to a management VRF instance into another VRF instance.

Use the `no` form of this command to remove the imported routes.

Example

```
OS10(conf-vrf)# ip route-import 1:1 ==> No route-map attached
```

Supported Releases 10.4.3.0 or later

ip route-export

Exports an IPv4 static route from one VRF instance to another.

Syntax

```
[no] ip route-export route-target [route-map route-map-name]
```

Parameters

- *route-target* — Enter the `route-target` of the VRF instance.
- `route-map` *route-map-name* — (Optional) Enter the route-map name to specify the route-map.

Default Not configured

Command Mode VRF CONFIG

Usage Information You can export routes corresponding only to a non-default or a default VRF instance. You cannot export routes that belong to a management VRF instance.

Use the `no` form of this command to undo the configuration.

Example

```
OS10(conf-vrf)# ip route-export 1:1 ==> No route-map attached
OS10(conf-vrf)# ip route-export 1:1 route-map abc ==> Route-map abc attached
to filter export routes
```

Supported Releases 10.4.3.0 or later

ipv6 route-import

Imports an IPv6 static route into a VRF instance from another VRF instance.

Syntax

```
[no] ipv6 route-import route-target
```

Parameters

- *route-target* — Enter the `route-target` of the VRF instance.

Default Not configured

Command Mode VRF CONFIG

Usage Information You can import IPv6 routes corresponding only to a non-default or a default VRF instance. You cannot import IPv6 routes that belong to a management VRF instance into another VRF instance.

Use the `no` form of this command to remove the imported routes.

Example

```
OS10(conf-vrf)# ipv6 route-import 1:1 ==> No route-map attached
```

Supported Releases 10.4.3.0 or later

ipv6 route-export

Exports an IPv6 static route from a VRF instance to another VRF instance.

Syntax `[no] ipv6 route-export route-target [route-map route-map-name]`

Parameters

- *route-target* — Enter the `route-target` of the VRF instance.
- `route-map` *route-map-name* — (Optional) Enter the route-map name to specify the route-map.

Default Not configured

Command Mode VRF CONFIG

Usage Information You can export IPv6 routes corresponding only to a non-default or a default VRF instance. You cannot export IPv6 routes that belong to a management VRF instance into another VRF instance.

Use the `no` form of this command to undo the configuration.

Example

```
OS10(conf-vrf)# ipv6 route-export 1:1 ==> No route-map attached
OS10(conf-vrf)# ipv6 route-export 1:1 route-map abc ==> Route-map abc attached
to filter export routes
```

Supported Releases 10.4.3.0 or later

ip scp vrf

Configures an SCP connection for the management or non-default VRF instance.

Syntax `ip scp vrf {management | vrf vrf-name}`

Parameters

- `management` — Enter the keyword to configure an SCP connection for the management VRF instance.
- `vrf` *vrf-name* — Enter the keyword then the name of the VRF to configure an SCP connection for that VRF instance.

Default Not configured

Command Mode CONFIGURATION

Usage Information	The `no` version of this command removes the management VRF instance configuration from the SCP client.
Example	`OS10(config)# ip scp vrf management` `OS10(config)# ip scp vrf vrf-blue`
Supported Releases	10.4.0E(R1) or later

ip sftp vrf

Configures an SFTP client for the management or non-default VRF instance.

Syntax	`ip sftp vrf {management | vrf vrf-name}`
Parameters	• `management` — Enter the keyword to configure an SFTP client for a management VRF instance. • `vrf vrf-name` — Enter the keyword then the name of the VRF to configure an SFTP client for that non-default VRF instance.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command removes the management VRF instance configuration from the SFTP client.
Example	`OS10(config)# ip sftp vrf management` `OS10(config)# ip sftp vrf vrf-blue`
Supported Releases	10.4.0E(R1) or later

ip tftp vrf

Configures a TFTP client for the management or non-default VRF instance.

Syntax	`ip tftp vrf {management | vrf vrf-name}`
Parameters	• `management` — Enter the keyword to configure a TFTP client for the management VRF instance. • `vrf vrf-name` — Enter the keyword then the name of the VRF to configure a TFTP client for that non-default VRF instance.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command removes the management VRF instance configuration from the TFTP client.
Example	`OS10(config)# ip tftp vrf management` `OS10(config)# ip tftp vrf vrf-blue`
Supported Releases	10.4.0E(R1) or later

ip vrf management

Configures the management VRF instance.

Syntax	`ip vrf management`

Parameters None

Default Not configured

Command Mode CONFIGURATION

Usage Information Enter the `ip vrf management` command only in non-transaction-based configuration mode. Do not use transaction-based mode. The `no` version of this command removes the management VRF instance configuration.

Example

```
OS10(config)# ip vrf management
OS10(conf-vrf)#
```

Supported Releases 10.4.0E(R1) or later

show hosts vrf

Displays the host table in the management or non-default VRF instance.

Syntax `show hosts vrf {management | vrf-name}`

Parameters

- `management`—Enter the keyword management to display the host table in the management VRF instance.
- *`vrf-name`*—Enter the name of the non-default VRF instance to display the host table in that VRF instance.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show hosts vrf management
Default Domain Name : dell.com
Domain List : abc.com xyz.net
Name Servers : 10.16.126.1
================================================
       Static Host to IP mapping Table
================================================
Host                                  IP-Address
------------------------------------------------
google.com                       172.217.160.142
yahoo.com                         98.139.180.180
```

Supported Releases 10.4.0E(R1) or later

show ip vrf

Displays the VRF instance information.

Syntax `show ip vrf [management | vrf-name]`

Parameters

- `management`—Enter the keyword `management` to display information corresponding to the management VRF instance.
- *`vrf-name`*—Enter the name of the non-default VRF instance to display information corresponding to that VRF instance.

Default Not configured

Command Mode EXEC

Usage Information	None

Example	`OS10# show ip vrf`
	`VRF-Name Interfaces`
	`default Mgmt1/1/1`
	`Eth1/1/1-1/1/2`
	`Vlan1`
	
	`management`
	
	`OS10# show ip vrf management`
	`VRF-Name Interfaces`
	`management`

Supported Releases	10.4.0E(R1) or later

update-source-ip

Configures a source IP interface for any leaked route in a VRF instance.

Syntax	`update-source-ip interface` *interface-id* To undo this configuration, use the `no update-source-ip` command.
Parameters	• `interface` *interface-id* — Enter the loopback interface identifier. The range is from 0 to 16383.
Default	Not configured
Command Mode	VRF CONFIGURATION
Example	`OS10(conf-vrf)# update-source-ip loopback 1`
Supported Releases	10.4.2E or later.

Bidirectional Forwarding Detection

The Bidirectional Forwarding Detection (BFD) protocol rapidly detects communication failures between two adjacent routers. BFD replaces link-state detection mechanisms in existing routing protocols. It also provides a failure detection solution for links with no routing protocols.

BFD provides forwarding-path failure detection in milliseconds instead of seconds. Because BFD is independent of routing protocols, it provides consistent network failure detection. BFD eliminates multiple protocol-dependent timers and methods. Networks converge is faster because BFD triggers link-state changes in the routing protocol sooner and more consistently.

BFD is a simple hello mechanism. Two neighboring routers running BFD establish a session using a three-way handshake. After the session is established, the routers exchange periodic control packets at sub-second intervals. If a router does not receive a hello packet within the specified time, routing protocols are notified that the forwarding path is down.

In addition, BFD sends a control packet when there is a state change or change in a session parameter. These control packets are sent without regard to transmit and receive intervals in a routing protocol.

BFD is an independent and generic protocol, which all media, topologies, and routing protocols can support using any encapsulation. OS10 implements BFD at Layer 3 (L3) and with User Datagram Protocol (UDP) encapsulation. BFD is supported on static and dynamic routing protocols, such as VRRP, OSPF, OSPFv3, IS-IS, and BGP.

The system displays BFD state change notifications.

NOTE: In this release, BFD is only supported for the border gateway protocol (BGP).

BFD session states

To establish a BFD session between two routers, enable BFD on both sides of the link. BFD routers can operate in both active and passive roles.

- The active router starts the BFD session. Both routers can be active in the same session.
- The passive router does not start a session. It only responds to a request for session initialization from the active router.

A BFD session can occur in Asynchronous and Demand modes. However, OS10 BFD supports only Asynchronous mode.

- In Asynchronous mode, both systems send periodic control messages at a specified interval to indicate that their session status is `Up`.
- In Demand mode, if one router requests Demand mode, the other router stops sending periodic control packets; it only sends a response to status inquiries from the Demand mode initiator. Either peer router, but not both, can request Demand mode at any time.

A BFD session can have four states: `Administratively Down`, `Down`, `Init`, and `Up`. The default BFD session state is `Down`.

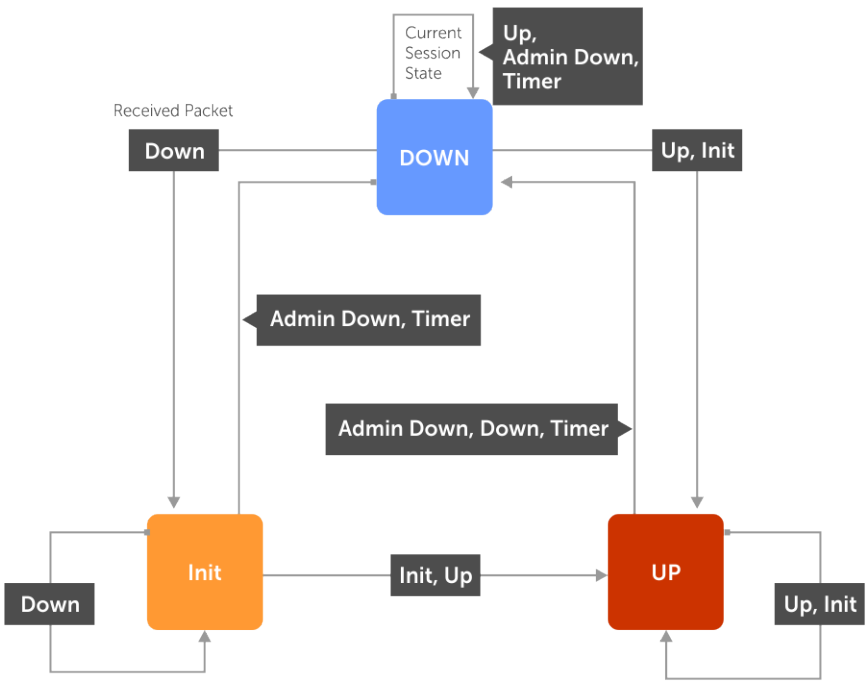
- `Administratively Down` — The local BFD router does not participate in the session.
- `Down` — The remote BFD router is not sending control packets or does not send them within the detection time for the session.
- `Init` — The local BFD outer is communicating to the remote router in the session.
- `Up` — Both BFD routers are sending control packets.

A BFD session's state changes to `Down` if:

- A control packet is not received within the detection time.
- Demand mode is active and a control packet is not received in response to a poll packet.

BFD session state changes example

The session state on a router changes according to the status notification it receives from the peer router. For example, if the current session state is `Down` and the router receives a `Down` status notification from the remote router, the session state on the local router changes to `Init`.



BFD three-way handshake

A BFD session requires a three-way handshake between neighboring routers. In the following example, the handshake assumes:

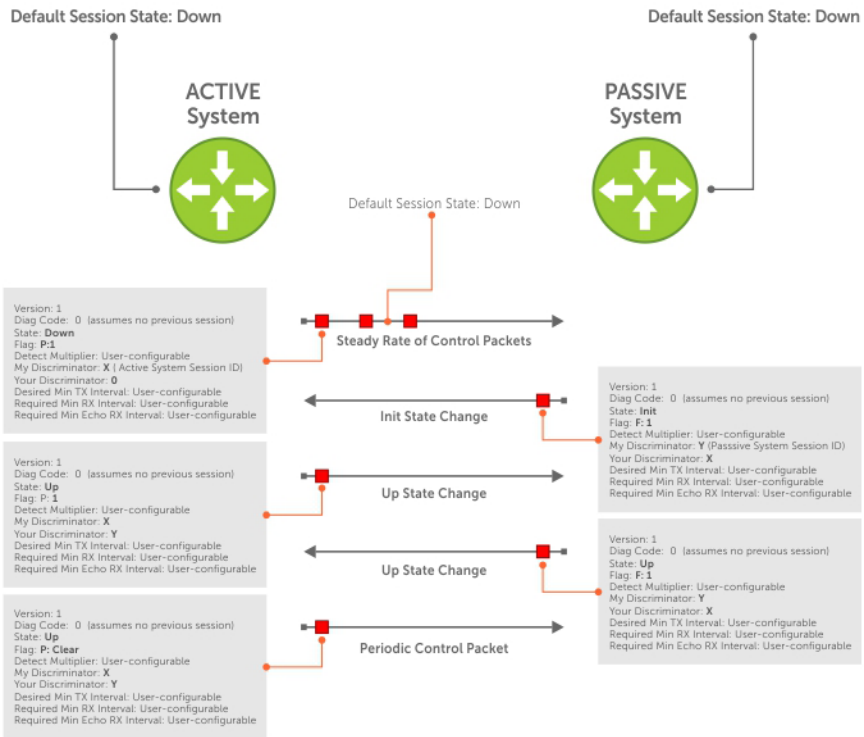
- One router is active, and the other router is passive.
- This is the first session established on this link.
- The default session state on both ports is `Down`.

1 The active system sends a steady stream of control packets to indicate that its session state is `Down` until the passive system responds. These packets are sent at the desired transmit interval of the Active system. The `Your Discriminator` field is set to zero.

2 When the passive system receives a control packet, it changes its session state to `Init` and sends a response to indicate its state change. The response includes its session ID in the `My Discriminator` field and the session ID of the remote system in the `Your Discriminator` field.

3 The active system receives the response from the passive system and changes its session state to `Up`. It then sends a control packet to indicate this state change. Discriminator values exchange, and transmit intervals negotiate.

4 The passive system receives the control packet and changes its state to `Up`. Both systems agree that a session is established. However, because both members must send a control packet, which requires a response, whenever there is a state change or change in a session parameter, the passive system sends a final response indicating the state change. After this, periodic control packets exchange.



BFD configuration

Before you configure BFD for a routing protocol, first enable BFD globally on both routers in the link. BFD is disabled by default.

- OS10 supports:
  - 64 BFD sessions at 100 minimum transmit and receive intervals with a multiplier of 4
  - 100 BFD sessions at 200 minimum transmit and receive intervals with a multiplier of 3
- OS10 does not support Demand mode, authentication, and Echo function.
- OS10 does not support BFD on multi-hop and virtual links.
- OS10 supports protocol liveness only for routing protocols.
- OS10 BFD supports only the BGP routing protocol. For IPv4 and IPv6 BGP, OS10 supports only the default virtual routing and forwarding (VRF).

Configure BFD globally

Before you configure BFD for static routing or a routing protocol, configure BFD globally on each router, including the global BFD session settings. BFD is disabled by default.

1 Configure the global BFD session parameters in CONFIGURATION mode.

```
bfd interval milliseconds min_rx milliseconds multiplier number role {active | passive}
```

- `interval milliseconds` — Enter the time interval for sending control packets to BFD peers, from 100 to 1000; default 200. Dell EMC recommends using more than 100 milliseconds.
- `min_rx milliseconds` — Enter the maximum waiting time for receiving control packets from BFD peers, from 100 to 1000; default 200. Dell EMC recommends using more than 100 milliseconds.
- `multiplier number` — Enter the number of consecutive packets that must not be received from a BFD peer before the session state changes to `Down`, from 3 to 50; default 3.
- `role {active | passive}` — Enter `active` if the router initiates BFD sessions. Both BFD peers can be active at the same time. Enter `passive` if the router does not initiate BFD sessions, and only responds to a request from an active BFD to initialize a session. The default is `active`.

2 Enable BFD globally in CONFIGURATION mode.

```
bfd enable
```

To verify that BFD is globally enabled, enter the `show running-config bfd` command.

BFD global configuration

```
OS10(config)# bfd interval 250 min_rx 300 multiplier 4 role passive
OS10(config)# bfd enable
OS10(config)# do show running-config bfd
!
bfd enable
bfd interval 250 min_rx 300 multiplier 4 role passive
```

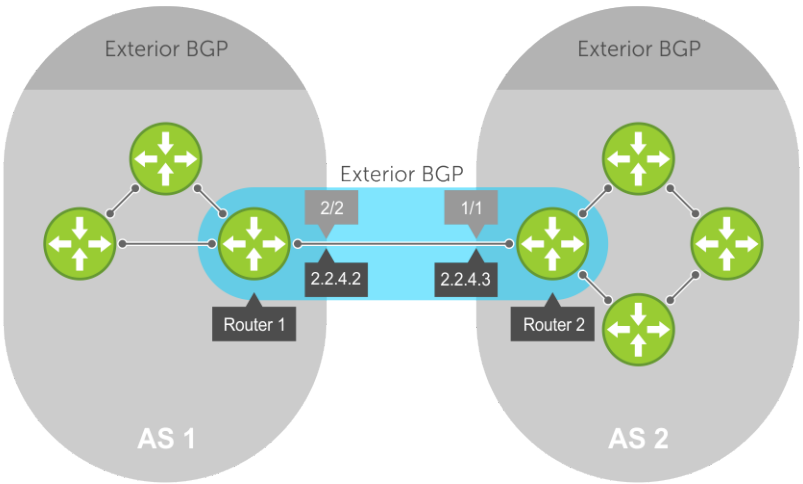
BFD for BGP

In a BGP core network, BFD enables faster network reconvergence. BFD rapidly detects communication failures in BGP fast-forwarding paths between internal BGP (iBGP) and external BGP (eBGP) peers. BFD for BGP is supported on physical, port-channel, and VLAN interfaces. BFD for BGP does not support the BGP multihop feature.

Before configuring BFD for BGP, first configure BGP on the interconnecting routers. For more information, see Border Gateway Protocol.

BFD for BGP example

In this BFD for BGP configuration example, Router 1 and Router 2 use eBGP in a transit network to interconnect AS1 and AS2. The eBGP routers exchange information with each other and with iBGP routers to maintain connectivity and accessibility within each autonomous system.



When you configure a BFD session with a BGP neighbor, you can:

- Establish a BFD session with a specified BGP neighbor using the `neighbor` *`ip-address`* and `bfd` commands.
- Establish BFD sessions with all neighbors discovered by BGP using the `bfd all-neighbors` command.

For example:

Router 1

```
OS10(conf)# bfd enable
OS10(conf)# router bgp 1
OS10(config-router-bgp-1)# neighbor 2.2.4.3
OS10(config-router-neighbor)# bfd
OS10(config-router-neighbor)# no shutdown
OR
OS10(conf)# bfd enable
OS10(conf)# router bgp 1
OS10(config-router-bgp-1)# bfd all-neighbors interval 200 min_rx 200 multiplier 6 role active
```

Router 2

```
OS10(conf)# bfd enable
OS10(conf)# router bgp 2
OS10(config-router-bgp-2)# neighbor 2.2.4.2
OS10(config-router-neighbor)# bfd
OS10(config-router-neighbor)# no shutdown
OR
OS10(conf)# bfd enable
OS10(conf)# router bgp 2
OS10(conf-router-bgp-2)# bfd all-neighbors interval 200 min_rx 200 multiplier 6 role active
```

BFD packets originating from a router are assigned to the highest priority egress queue to minimize transmission delays. Incoming BFD control packets received from the BGP neighbor are assigned to the highest priority queue within the control plane policing (CoPP) framework to avoid BFD packets drops due to queue congestion.

BFD notifies BGP of any failure conditions that it detects on the link. BGP initiates recovery actions.

BFD for BGP is supported only on directly connected BGP neighbors and in both BGP IPv4 and IPV6 networks. A maximum of 100 simultaneous BFD sessions are supported.

If each BFD for BGP neighbor receives a BFD control packet within the configured BFD interval for failure detection, the BFD session remains up and BGP maintains its adjacencies. If a BFD for BGP neighbor does not receive a control packet within the detection interval, the router informs any clients of the BFD session, and other routing protocols, about the failure. It then depends on the routing protocol that uses the BGP link to determine the appropriate response to the failure condition. The normal response is to terminate the peering session for the routing protocol and reconverge by bypassing the failed neighboring router. A log message generates whenever BFD detects a failure condition.

Configure BFD for BGP

OS10 supports BFD sessions with IPv4 or IPv6 BGP neighbors using the default VRF. When you configure BFD for BGP, you can enable BFD sessions with all BGP neighbors discovered by BGP or with a specified neighbor.

1 Configure BFD session parameters and enable BFD globally on all interfaces in CONFIGURATION mode as described in Configure BFD globally.

```
bfd interval milliseconds min_rx milliseconds multiplier number role {active | passive}
bfd enable
```

2 Enter the AS number of a remote BFD peer in CONFIGURATION mode, from 1 to 65535 for a 2-byte AS number and from 1 to 4294967295 for a 4-byte AS number. Only one AS number is supported per system. If you enter a 4-byte AS number, 4-byte AS support enables automatically.

```
router bgp as-number
```

3 Enter the IP address of a BFD peer in ROUTER-BGP mode. Enable a BFD session and the BGP link in ROUTER-NEIGHBOR mode. The global BFD session parameters configured in Step 1 are used.

```
neighbor ip-address
  bfd
  no shutdown
```

OR

Configure BFD sessions with all neighbors discovered by the BGP in ROUTER-BGP mode. The BFD session parameters you configure override the global session parameters configured in Step 1.

```
bfd all-neighbors [interval milliseconds min_rx milliseconds multiplier number role {active
| passive}]
```

- `interval milliseconds` — Enter the time interval for sending control packets to BFD peers, from 100 to 1000; default 200. Dell EMC recommends using more than 100 milliseconds.
- `min_rx milliseconds` — Enter the maximum waiting time for receiving control packets from BFD peers, from 100 to 1000; default 200. Dell EMC recommends using more than 100 milliseconds.
- `multiplier number` — Enter the maximum number of consecutive packets that are not received from a BFD peer before the session state changes to `Down`, from 3 to 50; default 3.
- `role {active | passive}` — Enter `active` if the router initiates BFD sessions. Both BFD peers can be active at the same time. Enter `passive` if the router does not initiate BFD sessions, and only responds to a request from an active BFD to initialize a session. The default is `active`.

To ignore the configured `bfd all-neighbors` settings for a specified neighbor, enter the `bfd disable` command in ROUTER-NEIGHBOR mode.

OR

Enter a BGP template with neighborhood name in ROUTER-BGP mode. Configure BFD sessions with all neighbors which inherit the template in ROUTER-TEMPLATE mode. For more information on how to use BGP templates, see Peer templates. The global BFD session parameters configured in Step 1 are used.

```
template template-name
  bfd
  no shutdown
```

4 Verify the BFD for BGP configuration in EXEC mode.

```
show bfd neighbors [detail]
```

BFD for BGP all-neighbors configuration

```
OS10(conf)# bfd interval 200 min_rx 200 multiplier 6 role active
OS10(conf)# bfd enable
OS10(conf)# router bgp 4
OS10(config-router-bgp-4)# bfd all-neighbors interval 200 min_rx 200 multiplier 6 role active
```

BFD for BGP single-neighbor configuration

```
OS10(conf)# bfd interval 200 min_rx 200 multiplier 6 role active
OS10(conf)# bfd enable
OS10(conf)# router bgp 1
OS10(config-router-bgp-1)# neighbor 150.150.1.1
OS10(config-router-neighbor)# bfd
OS10(config-router-neighbor)# no shutdown
```

BFD for BGP template configuration

```
OS10(config)# router bgp 300
OS10(config-router-bgp-300)# template ebgppg
OS10(config-router-template)# bfd
OS10(config-router-template)# exit
OS10(config-router-bgp-300)# neighbor 3.1.1.1
OS10(config-router-neighbor)# inherit template ebgppg
OS10(config-router-neighbor)# no shutdown
```

Display BFD operation

```
OS10# show bfd neighbors
* - Active session role
------------------------------------------------------------------------------
LocalAddr      RemoteAddr   Interface  State Rx-int Tx-int Mult VRF  Clients
------------------------------------------------------------------------------
* 150.150.1.2  150.150.1.1  vlan10     up    1000   1000   5   default  bgp
```

```
OS10# show bfd neighbors detail
Session Discriminator: 1
Neighbor Discriminator: 2
Local Addr: 150.150.1.2
Local MAC Addr: 90:b1:1c:f4:ab:fd
Remote Addr: 150.150.1.1
Remote MAC Addr: 90:b1:1c:f4:a4:d4
Interface: vlan10
State: up
Configured parameters:
TX: 1000ms, RX: 1000ms, Multiplier: 5
Actual parameters:
TX: 1000ms, RX: 1000ms, Multiplier: 5
Neighbor parameters:
TX: 200ms, RX: 200ms, Multiplier: 49
Role: active
VRF: default
Client Registered: bgp
Uptime: 01:58:09
Statistics:
  Number of packets received from neighbor: 7138
  Number of packets sent to neighbor: 7138
```

Verify BFD for BGP

```
OS10(config-router-bgp-101)# show ip bgp summary
BGP router identifier 30.1.1.2 local AS number 101
Global BFD is enabled
Neighbor    AS    MsgRcvd    MsgSent    Up/Down    State/Pfx
20.1.1.1    101   781        777        11:16:13     0
30.1.1.1    101   787        779        11:15:35     0
```

```
OS10(config-router-bgp-101)# show ip bgp neighbors
BGP neighbor is 20.1.1.1, remote AS 101, local AS 101  internal link

  BGP version 4, remote router ID 30.1.1.1
  BGP state ESTABLISHED, in this state for 11:19:01
  Last read 00:24:31 seconds
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time is 180, keepalive interval is 60 seconds
  Fall-over disabled
  Neighbor is using Global level BFD Configuration

  Received 784 messages
    1 opens, 0 notifications, 0 updates
    783 keepalives, 0 route refresh requests
  Sent 780 messages
    2 opens, 0 notifications, 0 updates
    778 keepalives, 0 route refresh requests
  Minimum time between advertisement runs is 30 seconds
  Minimum time before advertisements start is 0 seconds
  Capabilities received from neighbor for IPv4 Unicast:
    MULTIPROTO_EXT(1)
    ROUTE_REFRESH(2)
    CISCO_ROUTE_REFRESH(128)
    4_OCTET_AS(65)
  Capabilities advertised to neighbor for IPv4 Unicast:
    MULTIPROTO_EXT(1)
    ROUTE_REFRESH(2)
    CISCO_ROUTE_REFRESH(128)
    4_OCTET_AS(65)
  Prefixes accepted 0, Prefixes advertised 0
  Connections established 1; dropped 0
  Last reset never
  For address family: IPv4 Unicast
    Allow local AS number 0 times in AS-PATH attribute
   Prefixes ignored due to:
    Martian address  0, Our own AS in AS-PATH 0
    Invalid Nexthop  0, Invalid AS-PATH length 0
    Wellknown community 0, Locally originated 0

  Local host: 20.1.1.2, Local port: 179
  Foreign host: 20.1.1.1, Foreign port: 58248
```

BFD for OSPF

You can configure BFD to monitor and notify reachability status between OSPF neighbors. When you use BFD with OSPF, BFD sessions are established between all neighboring interfaces participating with OSPF full state. If a neighboring interface fails, BFD notifies OSPF protocol that a link state change has occurred.

To configure BFD for OSPF:

1 Enable BFD Globally.

2 Configure OSPF on the interconnecting routers. For more information, see Open Shortest Path First (OSPFv2 and OSPFv3).

Enable BFD Globally

To enable BFD globally:

Enable BFD globally.

```
bfd enable
```

CONFIGURATION Mode

Establishing BFD sessions with OSPFv2 neighbors

You can establish BFD sessions with all OSPF neighbors at one go. Alternatively, you can also establish BFD sessions with OSPF neighbors corresponding to a single OSPF interface.
To establish BFD sessions with OSPFv2 neighbors:

1 Enable BFD globally

```
bfd enable
```

CONFIGURATION Mode

2 Enter ROUTER-OSPF mode

```
router ospf ospf-instance
```

CONFIGURATION Mode

3 Establish sessions with all OSPFv2 neighbors.

```
bfd all-neighbors
```

ROUTER-OSPF Mode

4 Enter INTERFAC E CONFIGURATION mode.

```
interface interface-name
```

CONFIGURATION Mode

5 Establish BFD sessions with OSPFv2 neighbors corresponding to a single OSPF interface.

```
ip ospf bfd all-neighbors
```

INTERFACE CONFIGURATION Mode

Establishing BFD sessions with OSPFv2 neighbors in a non-default VRF instance

To establish BFD sessions with OSPFv2 neighbors in a non-default VRF instance:

1 Enable BFD globally

```
bfd enable
```

CONFIGURATION Mode

2 Enter INTERFACE CONFIGURATION mode

```
interface interface-name
```

CONFIGURATION Mode

3 Associate a non-default VRF with the interface you have entered.

```
ip vrf forwarding vrf1
```

INTERFACE CONFIGURATION Mode

4 Assign an IP address to the VRF.

```
ip address ip-address
```

VRF CONFIGURATION Mode

5 Attach the interface to an OSPF area.

```
ip ospf ospf-instance area area-address
```

VRF CONFIGURATION Mode

6 Establish BFD session with OSPFv2 neighbors in a single OSPF interface in a non-default VRF instance.

```
ip ospf bfd all-neoghbors
```

VRF CONFIGURATION Mode

7 Enter ROUTER-OSPF mode in a non-default VRF instance.

```
router ospf ospf-instance vrf vrf-name
```

8 Establish BFD sessions with all OSPFv2 instances in a non-default VRF.

```
bfd all-neighbors
```

```
OS10# show running-configuration ospf
!
interface vlan200
 no shutdown
 ip vrf forwarding red
 ip address 20.1.1.1/24
 ip ospf 200 area 0.0.0.0
 ip ospf bfd all-neighbors disable
!
interface vlan300
 no shutdown
 ip vrf forwarding red
 ip address 30.1.1.1/24
 ip ospf 200 area 0.0.0.0
!
router ospf 200 vrf red
 bfd all-neighbors
 log-adjacency-changes
 router-id 2.3.3.1
!
```

In this example OSPF is enabled in non-default VRF red. BFD is enabled globally at the router OSPF level and all the interfaces associated with this VRF OSPF instance inherit the global BFD configuration. However, this global BFD configuration does not apply to interfaces in which the interface level BFD configuration is already present. Also, VLAN 200 takes the interface level BFD configuration as interface-level BFD configuration takes precedent over the global OSPF-level BFD configuration.

Changing OSPFv2 BFD session parameters

Configure BFD sessions with default intervals and a default role.
The parameters that you can configure are: desired tx interval, required min rx interval, detection multiplier, and system role. Configure these parameters for all OSPF sessions or all OSPF sessions on a particular interface. If you change a parameter globally, the change affects all OSPF neighbors sessions. If you change a parameter at the interface level, the change affects all OSPF sessions on that interface.

NOTE: By default, OSPF uses the following BFD parameters for it's neighbors: min_tx = 200 msec, min_rx = 200 msec, multiplier = 3, role = active. If BFD is configured under interface context, that will be given high priority.

To change parameters for all OSPFv2 sessions or for OSPF sessions on a single interface, use the following commands:

1 Change parameters for OSPF sessions.

```
bfd all-neighbors interval milliseconds min_rx milliseconds multiplier value role [active | passive]
```

ROUTER-OSPF Mode

2 Change parameters for all OSPF sessions on an interface.

```
ip ospf bfd all-neighbors interval milliseconds min_rx milliseconds multiplier value role [active | passive]
```

INTERFACE CONFIGURATION Mode

Disabling BFD for OSPFv2

If you disable BFD globally, all sessions are torn down and sessions on the remote system are placed in a Down state. If you disable BFD on an interface, sessions on the interface are torn down and sessions on the remote system are placed in a Down state. Disabling BFD does not trigger a change in BFD clients; a final Admin Down packet is sent before the session is terminated.

To disable BFD sessions, use the following commands:

1 Disable BFD sessions with all OSPF neighbors.

```
no bfd all-neighbors
```

ROUTER-OSPF Mode

2 Disable OSPFv2 at interface level using the following command:

```
ip ospf bfd all-neighbors disable
```

INTERFACE CONFIGURATION Mode

To re-enable BFD, disabled the interface alone using the following commands:

- `no ip ospf bfd all-neighbors command`
- `ip ospf bfd all-neighbors`

Configure BFD for OSPFv3

BFD for OSPFv3 provides support for IPv6:

1 Enable BFD Globally.

2 Establish sessions with OSPFv3 neighbors.

Establishing BFD sessions with OSPFv3 neighbors

To establish BFD sessions with OSPFv3 neighbors:

1 Enable BFD globally

```
bfd enable
```

CONFIGURATION Mode

2 Enter ROUTER-OSPF mode

```
router ospfv3 ospfv3-instance
```

CONFIGURATION

3 Establish sessions with all OSPFv3 neighbors.

```
bfd all-neighbors
```

ROUTER-OSPFv3 Mode

4 Enter INTERFAC E CONFIGURATION mode.

```
interface interface-name
```

CONFIGURATION Mode

5 Establish BFD sessions with OSPFv3 neighbors corresponding to a single OSPF interface.

```
ipv6 ospf bfd all-neighbors
```

INTERFACE CONFIGURATION Mode

Establishing BFD sessions with OSPFv3 neighbors in a non-default VRF instance

To establish BFD sessions with OSPFv3 neighbors in a non-default VRF instance:

1 Enable BFD globally

```
bfd enable
```

CONFIGURATION Mode

2 Enter INTERFACE CONFIGURATION mode

```
interface interface-name
```

CONFIGURATION Mode

3 Associate a non-default VRF with the interface you have entered.

```
ip vrf forwarding vrf1
```

INTERFACE CONFIGURATION Mode

4 Assign an IP address to the VRF.

```
ip address ip-address
```

VRF CONFIGURATION Mode

5 Attach the interface to an OSPF area.

```
ipv6 ospf ospf-instance area area-address
```

VRF CONFIGURATION Mode

6 Establish BFD session with OSPFv3 neighbors in a single OSPF interface in a non-default VRF instance.

```
ipv6 ospf bfd all-neoghbors
```

VRF CONFIGURATION Mode

7 Enter ROUTER-OSPF mode in a non-default VRF instance.

```
router ospf ospf-instance vrf vrf-name
```

CONFIGURATION Mode

8 Establish BFD sessions with all OSPFv2 instances in a non-default VRF.

```
bfd all-neighbors
```

Changing OSPFv3 session parameters

Configure BFD sessions with default intervals and a default role.
The parameters that you can configure are: desired tx interval, required min rx interval, detection multiplier, and system role. Configure these parameters for all OSPFv3 sessions or all OSPFv3 sessions on a particular interface. If you change a parameter globally, the change affects all OSPFv3 neighbors sessions. If you change a parameter at the interface level, the change affects all OSPF sessions on that interface.

NOTE: By default, OSPF uses the following BFD parameters for it's neighbors: min_tx = 200 msec, min_rx = 200 msec, multiplier = 3, role = active. If BFD is configured under interface context, that will be given high priority.

To change parameters for all OSPFv3 sessions or for OSPF sessions on a single interface, use the following commands:

1 Change parameters for OSPF sessions.

```
bfd all-neighbors interval milliseconds min_rx milliseconds multiplier value role [active | passive]
```

ROUTER-OSPFv3 Mode

2 Change parameters for all OSPF sessions on an interface.

```
ipv6 ospf bfd all-neighbors interval milliseconds min_rx milliseconds multiplier value role [active | passive]
```

INTERFACE CONFIGURATION Mode

Disabling BFD for OSPFv3

If you disable BFD globally, all sessions are torn down and sessions on the remote system are placed in a Down state. If you disable BFD on an interface, sessions on the interface are torn down and sessions on the remote system are placed in a Down state. Disabling BFD does not trigger a change in BFD clients; a final Admin Down packet is sent before the session is terminated.

To disable BFD sessions, use the following commands:

1 Disable BFD sessions with all OSPF neighbors.

```
no bfd all-neighbors
```

ROUTER-OSPFv3 Mode

2 Disable BFD sessions with all OSPF neighbors on an interface.

```
ipv6 ospf bfd all-neighbors disable
```

INTERFACE CONFIGURATION Mode

To re-enable BFD, disabled the interface alone using the following commands:

- `no ipv6 ospf bfd all-neighbors command`
- `ipv6 ospf bfd all-neighbors`

BFD for Static route

The static Route BFD feature enables association of static routes with a BFD session in order to monitor the static route reachability. Depending on the status of the BFD session the static routes are added to or removed from the Routing Information Base (RIB). When BFD is configured, the nexthop reachability is dependent on the BFD state of the BFD session corresponding to the specified next hop. If the BFD session of the configured nexthop is down the static route will not be installed in the RIB.

The BFD session needs to be established successfully for the static route. BFD must be configured on both the peers pointing to its neighbor as next hop. There is no dependency on the order of configuration of static route and BFD configuration. The user has provision to configure BFD for all the static routes configured or for none of the static routes. Both IPv4 and IPv6 static route BFD is supported.

NOTE: You can configure BFD for all the static routes. Meaning, there is no provision for configuring BFD only for some of the existing static routes.

Configuring BFD for static routes is a three-step process:

1 Enable BFD Globally.

2 Configure static routes on both routers on the system (either local or remote). Configure static route in such a way that the next-hop interfaces point to each other.

3 Configure BFD for static route using the `ip route bfd` command

Establishing BFD Sessions for IPv4 Static Routes

Sessions are established for all neighbors that are the next hop of a static route.
To establish a BFD session, use the following command.

Establish BFD sessions for all neighbors that are the next hop of a static route.

```
ip route bfd [interval interval min_rx min_rx multiplier value role {active | passive}]
```

CONFIGURATION Mode

Establishing BFD Sessions for IPv4 Static Routes in a non-default VRF instance

To establish a BFD session for IPv4 static routes in a non-default VRF instance, use the following command.

Establish BFD sessions for all neighbors that are the next hop of a static route.

```
ip route bfd [vrf vrf-name] [interval interval min_rx min_rx multiplier value role {active | passive}]
```

CONFIGURATION Mode

Changing IPv4 Static Route Session Parameters

BFD sessions are configured with default intervals and a default role.
The parameters you can configure are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. These parameters are configured for all static routes. If you change a parameter, the change affects all sessions for static routes. To change parameters for static route sessions, use the following command.

Change parameters for all static route sessions.

```
ip route bfd interval milliseconds min_rx milliseconds multiplier value role [active | passive]
```

CONFIGURATION Mode

NOTE: By default, OSPF uses the following BFD parameters for it's neighbors: min_tx = 200 msec, min_rx = 200 msec, multiplier = 3, role = active. The values are configured in milliseconds

Disabling BFD for IPv4 Static Routes

If you disable BFD, all static route BFD sessions are torn down.
A final Admin Down packet is sent to all neighbors on the remote systems, and those neighbors change to the Down state. To disable BFD for IPv4 static routes, use the following command.

Disable BFD for static routes.

```
no ip route bfd
```

CONFIGURATION Mode

Establishing BFD Sessions for IPv6 Static Routes

To establish a BFD session for IPv6 static routes, use the following command.

Establish BFD sessions for all neighbors that are the next hop of a static route.

```
ipv6 route bfd [interval interval min_rx min_rx multiplier value role {active | passive}]
```

CONFIGURATION Mode

NOTE: By default, OSPF uses the following BFD parameters for it's neighbors: min_tx = 200 msec, min_rx = 200 msec, multiplier = 3, role = active. The values are configured in milliseconds

Establishing BFD Sessions for IPv6 Static Routes in a non-default VRF instance

To establish a BFD session for IPv6 static routes in a non-default VRF instance, use the following command.

Establish BFD sessions for all neighbors that are the next hop of a static route.

```
ipv6 route bfd [vrf vrf-name] [interval interval min_rx min_rx multiplier value role {active
| passive}]
```

CONFIGURATION Mode

NOTE: By default, OSPF uses the following BFD parameters for it's neighbors: min_tx = 200 msec, min_rx = 200 msec, multiplier = 3, role = active. The values are configured in milliseconds

Changing IPv6 Static Route Session Parameters

To change parameters for IPv6 static route sessions, use the following command.

Change parameters for all static route sessions.

```
ipv6 route bfd interval milliseconds min_rx milliseconds multiplier value role [active |
passive]
```

CONFIGURATION Mode

NOTE: By default, OSPF uses the following BFD parameters for it's neighbors: min_tx = 200 msec, min_rx = 200 msec, multiplier = 3, role = active. The values are configured in milliseconds

Disabling BFD for IPv6 Static Routes

To disable BFD for IPv6 static routes, use the following command.

Disable BFD for static routes.

```
no ipv6 route bfd
```

BFD commands

bfd

Enables BFD sessions with specified neighbors.

Syntax `bfd`

Parameters None

Default Not configured

Command Mode ROUTER-NEIGHBOR

ROUTER-TEMPLATE

Usage Information

- Use the `bfd` command to configure BFD sessions with a specified neighbor or neighbors which inherit a BGP template. Use the `neighbor {ip-address | ipv6-address}` command in ROUTER-BGP mode to specify the neighbor. Use the `template template-name` command in ROUTER-BGP mode to specify a BGP template. Use the `no bfd` command in ROUTER-NEIGHBOR mode to disable BFD sessions with a neighbor.
- Use the `bfd all-neighbors` command to configure L3 protocol-specific BFD parameters for all BFD sessions between discovered neighbors. The BFD parameters you configure override the global session parameters configured with the `bfd interval` command.

Example

```
OS10(conf)# router bgp 1
OS10(config-router-bgp-1)# neighbor 10.1.1.1
OS10(config-router-neighbor)# bfd
OS10(config-router-neighbor)# no shutdown
```

```
OS10(config)# router bgp 300
OS10(config-router-bgp-300)# template ebgppg
OS10(config-router-template)# bfd
OS10(config-router-template)# exit
OS10(config-router-bgp-300)# neighbor 3.1.1.1
OS10(config-router-neighbor)# inherit template ebgppg
OS10(config-router-neighbor)# no shutdown
```

Supported releases 10.4.1.0 or later

bfd all-neighbors

Configures all BFD session parameters established between neighbors discovered by an L3 protocol.

Syntax `bfd all-neighbors [milliseconds min_rx milliseconds multiplier number role {active | passive}]`

Parameters

- `interval milliseconds` — Enter the time interval for sending control packets to BFD peers, from 100 to 1000. Dell EMC recommends using more than 100 milliseconds.
- `min_rx milliseconds` — Enter the maximum waiting time for receiving control packets from BFD peers, from 100 to 1000. Dell EMC recommends using more than 100 milliseconds.

- `multiplier number` — Enter the maximum number of consecutive packets that must not be received from a BFD peer before the session state changes to `Down`, from 3 to 50.
- `role {active | passive}` — Enter `active` if the router initiates BFD sessions. Both BFD peers can be active at the same time. Enter `passive` if the router does not initiate BFD sessions, and only responds to a request from an active BFD to initialize a session.

Default The time interval for sending control packets to BFD peers is 200 milliseconds.

The maximum waiting time for receiving control packets from BFD peers is 200 milliseconds.

The number of consecutive packets that must be received from a BFD peer is 3.

The BFD role is `active`

Command Mode ROUTER-OSPF

Usage Information

- Use the `bfd all-neighbors` command to configure BFD sessions between discovered neighbors. The BFD session parameters you configure override the global session parameters configured with the `bfd interval` command. To disable BFD and ignore the configured `bfd all-neighbors` settings for a specified neighbor, enter the `bfd disable` command in ROUTER-NEIGHBOR mode.
- To remove the configured all-neighbors settings for all BGP neighbors, enter the `no` version of the command. To return to the default values, enter the `bfd all-neighbors` command.

Example

```
OS10(conf-ospf-bgp)# bfd all-neighbors interval 250 min_rx 300 multiplier 4
role passive
```

Supported releases 10.4.1.0 or later

bfd disable

Ignores the configured `bfd all-neighbors` settings and disables BFD for a specified neighbor.

Syntax `bfd disable`

Parameters None

Default Not configured

Command Mode ROUTER-NEIGHBOR

Usage Information Use the `neighbor ip-address` command in ROUTER-BGP mode to specify a neighbor. Use the `bfd disable` command to disable BFD sessions with the neighbor.

Example

```
OS10(conf)# router bgp 1
OS10(config-router-bgp-1)# neighbor 10.1.1.1
OS10(config-router-neighbor)# bfd disable
```

Supported releases 10.4.1.0 or later

bfd enable

Enables BFD on all interfaces on the switch.

Syntax `bfd enable`

Parameters None

Default	BFD is disabled.
Command Mode	CONFIGURATION
Usage Information	Before you configure BFD for static routing or a routing protocol, enable BFD globally on each router in a BFD session. To globally disable BFD on all interfaces, enter the `no bfd enable` command.
Example	`OS10(config)# bfd enable`
Supported releases	10.4.1.0 or later

bfd interval

Configures parameters for all BFD sessions on the switch.

Syntax `bfd interval *milliseconds* min_rx *milliseconds* multiplier *number* role {active | passive}`

Parameters

- `interval *milliseconds*` — Enter the time interval for sending control packets to BFD peers, from 100 to 1000. Dell EMC recommends using more than 100 milliseconds.
- `min_rx *milliseconds*` — Enter the maximum waiting time for receiving control packets from BFD peers, from 100 to 1000. Dell EMC recommends using more than 100 milliseconds.
- `multiplier *number*` — Enter the number of consecutive packets that must not be received from a BFD peer before the session state changes to `Down`, from 3 to 50.
- `role {active | passive}` — Enter `active` if the router initiates BFD sessions. Both BFD peers can be active at the same time. Enter `passive` if the router does not initiate BFD sessions, and only responds to a request from an active BFD to initialize a session.

Default

The time interval for sending control packets to BFD peers is 200 milliseconds.

The maximum waiting time for receiving control packets from BFD peers is 200 milliseconds.

The number of consecutive packets that must be received from a BFD peer is 3.

The BFD role is `active`.

Command Mode	CONFIGURATION
Usage Information	Use the `bfd interval` command to configure global BFD session settings. To configure the BFD parameters used in sessions established with neighbors discovered by an L3 protocol, use the `bfd all-neighbors` command. To remove the configured global settings and return to the default values, enter the `no` version of the command.
Example	`OS10(config)# bfd interval 250 min_rx 300 multiplier 4 role passive`
Supported releases	10.4.1.0 or later

ip ospf bfd all-neighbors

Enables and configures the default BFD parameters for all OSPFv2 neighbors in this interface.

Syntax `ip ospf bfd all-neighbors [disable|[interval *millisec* min_rx *min_rx* multiplier role {active | passive}]]`

Parameters

- `disable` — Disables the BFD session on an interface alone.
- `interval` *`milliseconds`* — Enter the time interval for sending control packets to BFD peers, from 100 to 1000. Dell EMC recommends using more than 100 milliseconds.
- `min_rx` *`milliseconds`* — Enter the maximum waiting time for receiving control packets from BFD peers, from 100 to 1000. Dell EMC recommends using more than 100 milliseconds.
- `multiplier` *`number`* — Enter the maximum number of consecutive packets that must not be received from a BFD peer before the session state changes to `Down`, from 3 to 50.
- `role {active | passive}` — Enter `active` if the router initiates BFD sessions. Both BFD peers can be active at the same time. Enter `passive` if the router does not initiate BFD sessions, and only responds to a request from an active BFD to initialize a session.

Default

The time interval for sending control packets to BFD peers is 200 milliseconds.

The maximum waiting time for receiving control packets from BFD peers is 200 milliseconds.

The number of consecutive packets that must be received from a BFD peer is 3.

The BFD role is `active`

Command Mode CONFIG-INTERFACE

Usage Information

- This command can be used to enable or disable BFD for an interface associated with OSPFv2. Interface level BFD configuration takes precedent over the OSPF global level BFD configuration. If there is no BFD configuration present at the interface level global OSPF BFD configuration will be inherited.

Example

```
(conf-if-eth1/1/1)#ip ospf bfd all-neighbors
```

Supported releases 10.4.2E or later

ipv6 ospf bfd all-neighbors

Enables and configures the default BFD parameters for all OSPFv3 neighbors in this interface.

Syntax

```
ipv6 ospf bfd all-neighbors [disable|[interval millisec min_rx min_rx
multiplier role {active | passive}]]
```

To disable default BFD parameters for all OSPFv3 neighbors using the `no ipv6 ospf bfd all-neighbors`.

Parameters

- `disable` — Disables the BFD session on an interface alone.
- `interval` *`milliseconds`* — Enter the time interval for sending control packets to BFD peers, from 100 to 1000. You cannot configure a value that is less than 100 milliseconds.
- `min_rx` *`milliseconds`* — Enter the maximum waiting time for receiving control packets from BFD peers, from 100 to 1000. Dell EMC recommends using more than 100 milliseconds.
- `multiplier` *`number`* — Enter the maximum number of consecutive packets that must not be received from a BFD peer before the session state changes to `Down`, from 3 to 50.
- `role {active | passive}` — Enter `active` if the router initiates BFD sessions. Both BFD peers can be active at the same time. Enter `passive` if the router does not initiate BFD sessions, and only responds to a request from an active BFD to initialize a session.

Default

The time interval for sending control packets to BFD peers is 200 milliseconds.

The maximum waiting time for receiving control packets from BFD peers is 200 milliseconds.

The number of consecutive packets that must be received from a BFD peer is 3.

The BFD role is `active`

Command Mode CONFIG-INTERFACE

Usage Information

- This command can be used to enable or disable BFD for an interface associated with OSPFv3. Interface level BFD configuration takes precedent over the OSPF global level BFD configuration. If there is no BFD configuration present at the interface level global OSPF BFD configuration will be inherited. All types of interfaces are supported.

Example

```
(conf-if-eth1/1/1)#ipv6 ospf bfd all-neighbors
```

Supported releases 10.4.2E or later

ip route bfd

Enables or disables BFD on static routes.

Syntax

```
ip route bfd [vrf vrf-name] [interval min_rx multiplier role {active | passive}]
```

To disable BFD on a static route, use the `no ip route bfd` command.

Parameters

- `vrf` *`vrf-name`* — Enter the keyword VRF followed by the name of the VRF to configure static route in that VRF.
- `interval` *`milliseconds`* — Enter the time interval for sending control packets to BFD peers, from 100 to 1000. Dell EMC recommends using more than 100 milliseconds.
- `min_rx` *`milliseconds`* — Enter the maximum waiting time for receiving control packets from BFD peers, from 100 to 1000. Dell EMC recommends using more than 100 milliseconds.
- `multiplier` *`number`* — Enter the maximum number of consecutive packets that must not be received from a BFD peer before the session state changes to `Down`, from 3 to 50.
- `role {active | passive}` — Enter `active` if the router initiates BFD sessions. Both BFD peers can be active at the same time. Enter `passive` if the router does not initiate BFD sessions, and only responds to a request from an active BFD to initialize a session.

Default The time interval for sending control packets to BFD peers is 200 milliseconds.

The maximum waiting time for receiving control packets from BFD peers is 200 milliseconds.

The number of consecutive packets that must be received from a BFD peer is 3.

The BFD role is `active`

Command Mode CONFIG

Usage Information

- This command can be used to enable or disable BFD for all the configured IPv4 static route for specified VRF. If VRF name is not specified the command will be applicable for default VRF.

Example

```
OS10(config)# ip route bfd interval 250 min_rx 250 multiplier 4 role active
```

Supported releases 10.4.2E or later

ipv6 route bfd

Enables or disables BFD on IPv6 static routes.

Syntax

```
ipv6 route bfd [vrf vrf-name] [interval millisec min_rx min_rx multiplier role {active | passive}]
```

To disable BFD on a IPv6 static route, use the `no ipv6 route bfd` command.

Parameters

- `vrf vrf-name` — Enter the keyword VRF followed by the name of the VRF to configure static route in that VRF.
- `interval milliseconds` — Enter the time interval for sending control packets to BFD peers, from 100 to 1000. You cannot configure a value that is less than 100 milliseconds.
- `min_rx milliseconds` — Enter the maximum waiting time for receiving control packets from BFD peers, from 100 to 1000. Dell EMC recommends using more than 100 milliseconds.
- `multiplier number` — Enter the maximum number of consecutive packets that must not be received from a BFD peer before the session state changes to `Down`, from 3 to 50.
- `role {active | passive}` — Enter `active` if the router initiates BFD sessions. Both BFD peers can be active at the same time. Enter `passive` if the router does not initiate BFD sessions, and only responds to a request from an active BFD to initialize a session.

Default

The time interval for sending control packets to BFD peers is 200 milliseconds.

The maximum waiting time for receiving control packets from BFD peers is 200 milliseconds.

The number of consecutive packets that must be received from a BFD peer is 3.

The BFD role is `active`

Command Mode CONFIG

Usage Information

- This command can be used to enable or disable BFD for all the configured IPv6 static route for specified VRF. If VRF name is not specified the command will be applicable for default VRF.

Example

```
OS10(config)# ipv6 route bfd interval 250 min_rx 250 multiplier 4 role active
```

Supported releases 10.4.2E or later

show bfd neighbors

Displays information about BFD neighbors from all interfaces using the default VRF.

Syntax

```
show bfd neighbors [detail]
```

Parameters `detail` — (Optional) View detailed information about BFD neighbors.

Default Not configured

Command Mode EXEC

Usage Information Use the `show bfd neighbors` command to verify that a BFD session between neighbors is up using the default VRF insta

Example

```
OS10# show bfd neighbors
* - Active session role
------------------------------------------------------------------------------
LocalAddr      RemoteAddr   Interface  State Rx-int Tx-int Mult VRF  Clients
------------------------------------------------------------------------------
* 150.150.1.2  150.150.1.1  vlan10     up    1000   1000   5   default  bgp
```

```
OS10# show bfd neighbors detail
Session Discriminator: 1
Neighbor Discriminator: 2
Local Addr: 150.150.1.2
Local MAC Addr: 90:b1:1c:f4:ab:fd
Remote Addr: 150.150.1.1
Remote MAC Addr: 90:b1:1c:f4:a4:d4
Interface: vlan10
State: up
Configured parameters:
TX: 1000ms, RX: 1000ms, Multiplier: 5
Actual parameters:
TX: 1000ms, RX: 1000ms, Multiplier: 5
Neighbor parameters:
TX: 200ms, RX: 200ms, Multiplier: 49
Role: active
VRF: default
Client Registered: bgp
Uptime: 01:58:09
Statistics:
  Number of packets received from neighbor: 7138
  Number of packets sent to neighbor: 7138
```

```
show bfd neighbors
* - Active session role
------------------------------------------------------------------------------
LocalAddr    RemoteAddr  Interface  State   RxInt   TxInt  Mult  VRF   Clients
------------------------------------------------------------------------------
100.1.3.2   100.1.3.1  vlan102     up     200      200   3    default ospfv2
* 100.1.4.2   100.1.4.1  vlan103     up     200      200   3    default ospfv2
* 100.1.5.2   100.1.5.1  vlan104     up     200      200   3    default ospfv2
* 100.1.6.2   100.1.6.1  vlan105     up     200      200   3    default ospfv2
* 100.1.7.2   100.1.7.1  vlan106     up     200      200   3    default ospfv2
* 100.1.8.2   100.1.8.1  vlan107     up     1000     1000  3    default ospfv2
* 100.1.9.2   100.1.9.1  vlan108     up     200      200   3    default ospfv2
```

Supported releases 10.4.1.0 or later

Border Gateway Protocol

Border Gateway Protocol (BGP) is an interautonomous system routing protocol that transmits interdomain routing information within and between autonomous systems (AS). BGP exchanges network reachability information with other BGP systems. BGP adds reliability to network connections by using multiple paths from one router to another. Unlike most routing protocols, BGP uses TCP as its transport protocol.

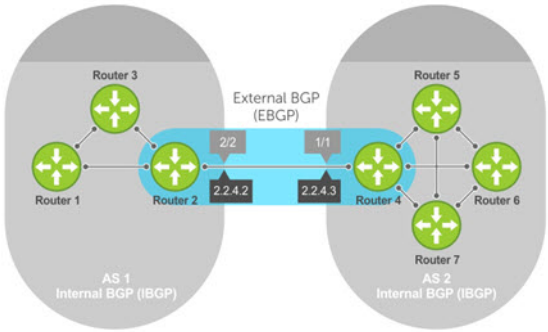
Autonomous systems

BGP autonomous systems are a collection of nodes under a single administration with shared network routing policies. Each AS has a number, which an Internet authority assigns—you do not assign the BGP number.

The Internet Assigned Numbers Authority (IANA) identifies each network with a unique AS number (ASN). AS numbers 64512 through 65534 are reserved for private purposes. AS numbers 0 and 65535 cannot be used in a live environment. IANA assigns valid AS numbers in the range of 1 to 64511.

Multihomed AS	Maintains connections to more than one other AS. This group allows the AS to remain connected to the Internet if a complete failure occurs to one of their connections. This type of AS does not allow traffic from one AS to pass through on its way to another AS.
Stub AS	Connected to only one AS.
Transit AS	Provides connections through itself to separate networks. For example, Router 1 uses Router 2—the transit AS, to connect to Router 4. Internet service providers (ISPs) are always a transit AS because they provide connections from one network to another. An ISP uses a transit AS to sell transit service to a customer network.

When BGP operates inside an AS - AS1 **or** AS2, it functions as an Internal Border Gateway Protocol (IBGP). When BGP operates between AS endpoints - AS1 **and** AS2, it functions as an External Border Gateway Protocol (EBGP). IBGP provides routers inside the AS with the path to reach a router external to the AS. EBGP routers exchange information with other EBGP routers and IBGP routers to maintain connectivity and accessibility.



Classless interdomain routing

BGPv4 supports classless interdomain routing (CIDR) with aggregate routes and AS paths. CIDR defines a network using a prefix consisting of an IP address and mask, resulting in efficient use of the IPv4 address space. Using aggregate routes reduces the size of routing tables.

Path-vector routing

BGP uses a path-vector protocol that maintains dynamically updated path information. Path information updates which returns to the originating node are detected and discarded. BGP does not use a traditional Internal Gateway Protocol (IGP) matrix but makes routing decisions based on path, network policies, and/or rule sets.

Full-mesh topology

In an AS, a BGP network must be in `full mesh` for routes received from an internal BGP peer to send to another IBGP peer. Each BGP router talks to all other BGP routers in a session. For example, in an AS with four BGP routers, each router has three peers; in an AS with six routers, each router has five peers.

Sessions and peers

A BGP session starts with two routers communicating using the BGP. The two end-points of the session are called *peers*. A peer is also called a *neighbor*. Events and timers determine the information exchange between peers. BGP focuses on traffic routing policies.

Sessions

In operations with other BGP peers, a BGP process uses a simple finite state machine consisting of six states—`Idle`, `Connect`, `Active`, `OpenSent`, `OpenConfirm`, and `Established`. For each peer-to-peer session, a BGP implementation tracks the state of the session. The BGP defines the messages that each peer exchanges to change the session from one state to another.

Idle	BGP initializes all resources, refuses all inbound BGP connection attempts, and starts a TCP connection to the peer.
Connect	Router waits for the TCP connection to complete and transitions to the `OpenSent` state if successful. If that transition is not successful, BGP resets the `ConnectRetry` timer and transitions to the `Active` state when the timer expires.
Active	Router resets the `ConnectRetry` timer to zero and returns to the `Connect` state.
OpenSent	Router sends an `Open` message and waits for one in return after a successful `OpenSent` transition.
OpenConfirm	Neighbor relation establishes and is in the `OpenConfirm` state after the Open message parameters are agreed on between peers. The router then receives and checks for agreement on the parameters of the open messages to establish a session.
Established	Keepalive messages exchange, and after a successful receipt, the router is in the `Established` state. Keepalive messages continue to send at regular periods. The keepalive timer establishes the state to verify connections.

After the connection is established, the router sends and receives keepalive, update, and notification messages to and from its peer.

Peer templates

Peer templates allow BGP neighbors to inherit the same outbound policies. Instead of manually configuring each neighbor with the same policy, you can create a peer group with a shared policy that applies to individual peers. A peer template provides efficient update calculation with a simplified configuration.

Peer templates also aid in convergence speed. When a BGP process sends the same information to many peers, a long output queue may be set up to distribute the information. For peers that are members of a peer template, the information is sent to one place then passed on to the peers within the template.

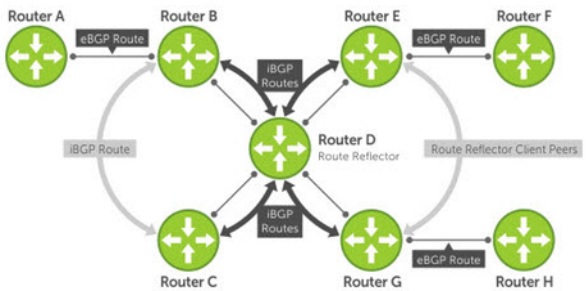
Route reflectors

Route reflectors (RRs) reorganize the IBGP core into a hierarchy and allow route advertisement rules. Route reflection divides IBGP peers into two groups — client peers and nonclient peers.

- If a route is received from a nonclient peer, it reflects the route to all client peers
- If a route is received from a client peer, it reflects the route to all nonclient and client peers

An RR and its client peers form a *route reflection cluster*. BGP speakers announce only the best route for a given prefix. RR rules apply after the router makes its best path decision.

NOTE: Do not use RRs in forwarding paths — hierarchal RRs that maintain forwarding plane RRs could create route loops.



Routers B, C, D, E, and G are members of the same AS—AS100. These routers are also in the same route reflection cluster, where Router D is the route reflector. Routers E and G are client peers of Router D, and Routers B and C and nonclient peers of Router D.

1 Router B receives an advertisement from Router A through EBGP. Because the route is learned through EBGP, Router B advertises it to all its IBGP peers — Routers C and D.

2 Router C receives the advertisement but does not advertise it to any peer because its only other peer is Router D (an IBGP peer) and Router D has already learned it through IBGP from Router B.

3 Router D does not advertise the route to Router C because Router C is a nonclient peer. The route advertisement came from Router B which is also a nonclient peer.

4 Router D does reflect the advertisement to Routers E and G because they are client peers of Router D.

5 Routers E and G advertise this IBGP learned route to their EBGP peers — Routers F and H.

Multiprotocol BGP

Multiprotocol BGP (MBGP) is an extension to BGP that supports multiple address families—IPv4 and IPv6. MBGP carries multiple sets of unicast and multicast routes depending on the address family.

You can enable the MBGP feature on a per router, per template, and/or a per peer basis. The default is the IPv4 unicast routes.

BGP session supports multiple address family interface (AFI) and sub address family interface (SAFI) combinations, BGP uses OPEN message to convey this information to the peers. As a result, the IPv6 routing information is exchanged over the IPv4 peers and vice versa.

BGP routers that support IPv6 can set up BGP sessions using IPv6 peers. If the existing BGP-v4 session is capable of exchanging ipv6 prefixes, the same is used to carry ipv4 as well as ipv6 prefixes. If the BGP-v4 neighbor goes down, it also impacts the IPv6 route exchange. If BGP-v6 session exists, it continues to operate independently from BGP-v4.

Multiprotocol BGPv6 supports many of the same features and functionality as BGPv4. IPv6 enhancements to MBGP include support for an IPv6 address family and Network Layer Reachability Information (NLRI) and next hop attributes that use the IPv6 addresses.

Attributes

Routes learned using BGP have associated properties that are used to determine the best route to a destination when multiple paths exist to a particular destination. These properties are called *BGP attributes* which influence route selection for designing robust networks. There are no hard-coded limits on the number of supported BGP attributes.

BGP attributes for route selection:

- Weight
- Local preference
- Multiexit discriminators
- Origin
- AS path
- Next-hop

Communities

BGP communities are sets of routes with one or more common attributes. Communities assign common attributes to multiple routes at the same time. Duplicate communities are not rejected.

Selection criteria

Best path selection criteria for BGP attributes:

1 Prefer the path with the largest WEIGHT attribute, and prefer the path with the largest LOCAL_PREF attribute.

2 Prefer the path that is locally originated using the `network` command, `redistribute` command, or `aggregate-address` command. Routes originated using a `network` or `redistribute` command are preferred over routes that originate with the `aggregate-address` command.

3 (Optional) If you configure the `bgp bestpath as-path ignore` command, skip this step because AS_PATH is not considered. Prefer the path with the shortest AS_PATH:
   - An AS_SET has a path length of 1 no matter how many are in the set
   - A path with no AS_PATH configured has a path length of 0
   - AS_CONFED_SET is not included in the AS_PATH length
   - AS_CONFED_SEQUENCE has a path length of 1 no matter how many ASs are in the AS_CONFED_SEQUENCE

4 Prefer the path with the lowest ORIGIN type—IGP is lower than EGP and EGP is lower than INCOMPLETE.

5 Prefer the path with the lowest multiexit discriminator (MED) attribute:
   - This comparison is only done if the first neighboring AS is the same in the two paths. The MEDs compare only if the first AS in the AS_SEQUENCE is the same for both paths.
   - Configure the `bgp always-compare-med` command to compare MEDs for all paths.
   - Paths with no MED are treated as "worst" and assigned a MED of 4294967295.

6 Prefer external (EBGP) to internal (IBGP) paths or confederation EBGP paths, and prefer the path with the lowest IGP metric to the BGP next-hop.

7 The system deems the paths as equal and only performs the following steps if the criteria are not met:
   - Configure the IBGP multipath or EBGP multipath using the `maximum-path` command.
   - The paths being compared were received from the same AS with the same number of AS in the AS Path but with different next-hops.
   - The paths were received from IBGP or EBGP neighbor, respectively.

8 If you enable the `bgp bestpath router-id ignore` command and:
   - If the Router-ID is the same for multiple paths because the routes were received from the same route—skip this step.
   - If the Router-ID is **not** the same for multiple paths, prefer the path that was first received as the Best Path. The path selection algorithm returns without performing any of the checks detailed.

9 Prefer the external path originated from the BGP router with the lowest router ID. If both paths are external, prefer the oldest path—first received path. For paths containing an RR attribute, the originator ID is substituted for the router ID. If two paths have the same router ID, prefer the path with the lowest cluster ID length. Paths without a cluster ID length are set to a 0 cluster ID length.

10 Prefer the path originated from the neighbor with the lowest address. The neighbor address is used in the BGP neighbor configuration and corresponds to the remote peer used in the TCP connection with the local router.

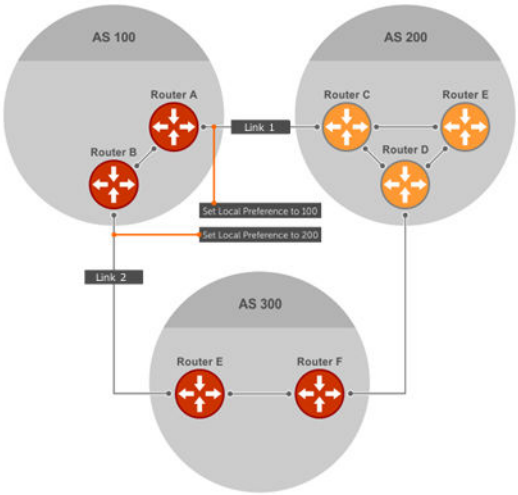
In Non-Deterministic mode, the `bgp non-deterministic-med` command applies. Paths compare in the order they arrive. This method leads to system selection of different best paths from a set of paths. Depending on the order they were received from the neighbors, MED may or may not get compared between the adjacent paths. In Deterministic mode, the system compares MED. MED is compared between the adjacent paths within an AS group because all paths in the AS group are from the same AS.

Weight and local preference

The weight attribute is local to the router and does not advertise to neighboring routers. If the router learns about more than one route to the same destination, the route with the highest weight is preferred. The route with the highest weight is installed in the IP routing table.

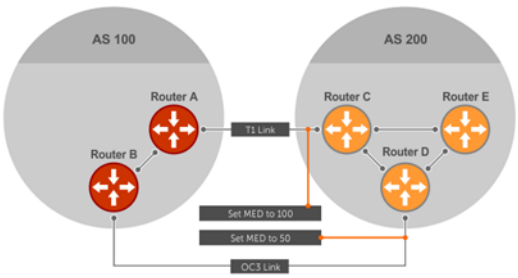
The local preference — LOCAL_PREF represents the degree of preference within the entire AS. The higher the number, the greater the preference for the route.

LOCAL_PREF is one of the criteria that determines the best path — other criteria may impact selection, see Best path selection. Assume that LOCAL_PREF is the only attribute applied and AS 100 has two possible paths to AS 200. Although the path through Router A is shorter, the LOCAL_PREF settings have the preferred path going through Router B and AS 300. This advertises to all routers within AS 100, causing all BGP speakers to prefer the path through Router B.



Multiexit discriminators

If two autonomous systems connect in more than one place, use a multiexit discriminator (MED) to assign a preference to a preferred path. MED is one of the criteria used to determine best path—other criteria may also impact selection.



One AS assigns the MED a value. Other AS uses that value to decide the preferred path. Assume that the MED is the only attribute applied and there are two connections between AS 100 and AS 200. Each connection is a BGP session. AS 200 sets the MED for its Link 1 exit point to 100 and the MED for its Link 2 exit point to 50. This sets up a path preference through Link 2. The MEDs advertise to AS 100 routers so they know which is the preferred path.

MEDs are nontransitive attributes. If AS 100 sends the MED to AS 200, AS 200 does not pass it on to AS 300 or AS 400. The MED is a locally relevant attribute to the two participating AS — AS 100 and AS 200. The MEDs advertise across both links—if a link goes down, AS 100 has connectivity to AS 300 and AS 400.

Origin

The origin indicates how the prefix came into BGP. There are three origin codes—IGP, EGP, and INCOMPLETE.

IGP	Prefix originated from information learned through an IGP.
EGP	Prefix originated from information learned from an EGP, which Next Generation Protocol (NGP) replaced.
INCOMPLETE	Prefix originated from an unknown source.

An IGP indicator means that the route was derived inside the originating AS. EGP means that a route was learned from an external gateway protocol. An INCOMPLETE origin code results from aggregation, redistribution, or other indirect ways of installing routes into BGP.

The question mark (?) indicates an origin code of INCOMPLETE, and the lower case letter (i) indicates an origin code of IGP.

Origin configuration

```
OS10# show ip bgp
BGP local RIB : Routes to be Added , Replaced , Withdrawn
BGP local router ID is 30.1.1.1
Status codes: s suppressed, S stale, d dampened, h history, * valid, > best
Path source: I - internal, a - aggregate, c - confed-external, r - redistributed
n - network S - stale
Origin codes: i - IGP, e - EGP, ? - incomplete
          Network       Next Hop  Metric LocPrf Weight Path
*>I       1.1.1.0/24    17.1.1.2  0      0      0      i
*>I       2.2.2.0/24    17.1.1.2  0      0      0      ?
*>I       3.3.3.0/24    17.1.1.2  0      0      0      e
```

AS path and next-hop

The AS path is the AS list that all the prefixes listed in the update have passed through. The BGP speaker adds the local AS number when advertising to an EBGP neighbor. Any update that contains the AS path number 0 is valid.

The next-hop is the IP address used to reach the advertising router:

- For EBGP neighbors, the next-hop address is the IP address of the connection between neighbors.
- For IBGP neighbors, the EBGP next-hop address is carried into the local AS. A next hop attribute sets when a BGP speaker advertises itself to another BGP speaker outside the local AS and when advertising routes within an AS.

For EBGP neighbors, the next-hop address corresponding to a BGP route does not resolve if the next-hop address is not the same as the neighbor IP address. The next-hop attribute also serves as a way to direct traffic to another BGP speaker, instead of waiting for a speaker to advertise. When a next-hop BGP neighbor is unreachable, the connection to that BGP neighbor goes down after the hold-down timer expires.

When you enable `fast-external-fallover` and if the router has learned the routes from the BGP neighbor, the BGP session terminates immediately if the next-hop becomes unreachable, without waiting for the hold-down time.

Best path selection

Best path selection selects the best route out of all paths available for each destination, and records each selected route in the IP routing table for traffic forwarding. Only valid routes are considered for best path selection. BGP compares all paths, in the order in which they arrive, and selects the best paths. Paths for active routes are grouped in ascending order according to their neighboring external AS number.

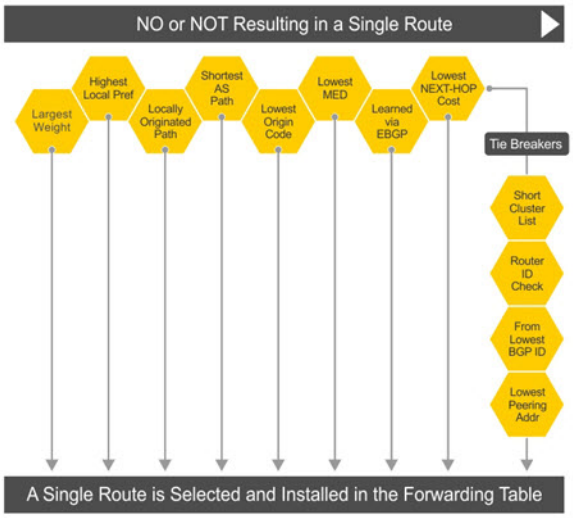
OS10 follows *deterministic* MED to select different best paths from a set of paths. This may depend on the order the different best paths are received from the neighbors — MED may or may not get compared between adjacent paths. BGP best path selection is deterministic by default.

The best path in each group is selected based on specific criteria—only one best path is selected at a time. If BGP receives more than one best path, it moves on to the next list of valid paths, and continues until it reaches the end of the list.

When you configure the `non-deterministic-med` command, paths are compared in the order they arrive. OS10 follows this method to select different best paths from a set of paths, depending on the order they were received from the neighbors—MED may or may not get compared between the adjacent paths.

By default, the `bestpath as-path multipath-relax` command is disabled. This prevents BGP from load-balancing a learned route across two or more EBGP peers. To enable load-balancing across different EBGP peers, enter the `bestpath as-path multipath-relax` command.

If you configure the `bgp bestpath as-path ignore` command and the `bestpath as-path multipath-relax` command at the same time, an error message displays—only enable one command at a time.



More path support

More path (Add-Path) reduces convergence times by advertising multiple paths to its peers for the same address prefix without replacing existing paths with new ones. By default, a BGP speaker advertises only the best path to its peers for a given address prefix.

If the best path becomes unavailable, the BGP speaker withdraws its path from its local router information base (RIB) and recalculates a new best path. This situation requires both IGP and BGP convergence and is a lengthy process. BGP add-path also helps switch over to the next new best path when the current best path is unavailable.

The Add-Path capability to advertise more paths is supported only on IBGP peers—it is not supported on EBGP peers or BGP peer groups.

Ignore router ID calculations

Avoid unnecessary BGP best path transitions between external paths under certain conditions. The `bestpath router-id ignore` command reduces network disruption caused by routing and forwarding plane changes and allows for faster convergence.

Advertise cost

As the default process for redistributed routes, OS10 supports IGP cost as MED. Both auto-summarization and synchronization are disabled by default.

BGPv4 and BGPv6 support

- Deterministic MED, default
- A path with a missing MED is treated as worst path and assigned an `0xffffffff` MED value
- Delayed configuration at system boot — OS10 reads the entire configuration file BEFORE sending messages to start BGP peer sessions

4-Byte AS numbers

OS10 supports 4-byte AS number configurations by default. The 4-byte support is advertised as a new BGP capability - `4-BYTE-AS`, in the OPEN message. A BGP speaker that advertises 4-Byte-AS capability to a peer, and receives the same from that peer must encode AS numbers as 4-octet entities in all messages.

If the AS number of the peer is different, the 4-byte speaker brings up the neighbor session using a reserved 2-byte ASN, 23456 called *AS_TRANS*. The AS_TRANS is used to interop between a 2-byte and 4-byte AS number.

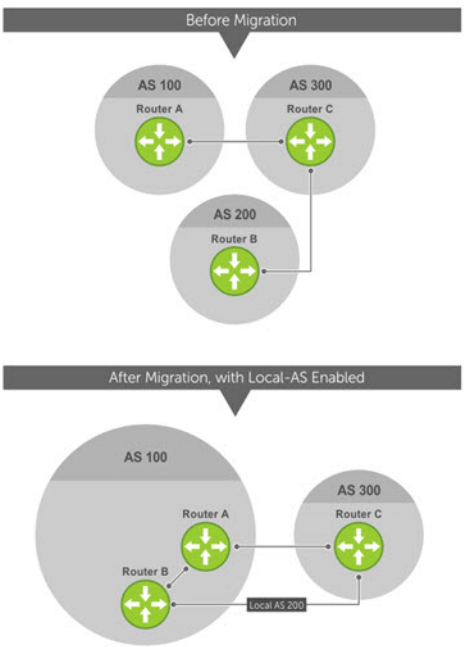
Where the 2-byte format is 1 to 65535, the 4-byte format is 1 to 4294967295. You can enter AS numbers using the traditional format.

AS number migration

You can transparently change the AS number of an entire BGP network. Changing the AS number ensures that the routes propagate throughout the network while migration is in progress. When migrating one AS to another and combining multiple AS, an EBGP network may lose its routing to an IBGP if the AS number changes.

Migration is difficult as all IBGP and EBGP peers of the migrating network must be updated to maintain network reachability. Local-AS allows the BGP speaker to operate as if it belongs to a virtual AS network besides its physical AS network.

Disable the `local-as` command after migration. Failure to disable the `local-as` command after migration causes the `local-as` command to replace the original AS number of the system. You must reconfigure the system with a new AS number.



Router A, Router B, and Router C belong to AS 100, 200, and 300, respectively. Router A acquired Router B — Router B has Router C as its client. When Router B is migrating to Router A, it must maintain the connection with Router C without immediately updating Router C's configuration. Local-AS allows Router B to appear as if it still belongs to Router B's old network, AS 200, to communicate with Router C.

The Local-AS does not prepend the updates with the AS number received from the EBGP peer if you use the `no prepend` command. If you do not select `no prepend`, the default, the Local-AS adds to the first AS segment in the AS-PATH. If you use an inbound route-map to prepend the AS-PATH to the update from the peer, the Local-AS adds first.

If Router B has an inbound route-map applied on Router C to prepend *65001 65002* to the AS-PATH, these events take place on Router B:

- Receive and validate the update.
- Prepend local-as 200 to AS-PATH.
- Prepend *65001 65002* to AS-PATH.

Local-AS prepends before the route map to give the appearance that the update passed through a router in AS 200 before it reaches Router B.

Configure Border Gateway Protocol

BGP is disabled by default. To enable the BGP process and start to exchange information, assign an AS number and use commands in ROUTER-BGP mode to configure a BGP neighbor.

BGP neighbor adjacency changes	All BGP neighbor changes are logged
Fast external fallover	Enabled
Graceful restart	Disabled
Local preference	100
4-byte AS	Enabled
MED	0
Route flap dampening parameters	• half-life = 15 minutes • max-suppress-time = 60 minutes • reuse = 750 • suppress = 2000
Timers	• keepalive = 60 seconds • holdtime = 180 seconds
Add-path	Disabled

Enable BGP

Before enabling BGP, assign a BGP router ID to the switch using the following command:

- In the ROUTER BGP mode, enter the `router-id` *`ip-address`* command. Where in, `ip-address` is the IP address corresponding to a configured L3 interface (physical, loopback, or LAG).

BGP is disabled by default. The system supports one AS number — you must assign an AS number to your device. To establish BGP sessions and route traffic, configure at least one BGP neighbor or peer. In BGP, routers with an established TCP connection are called

neighbors or *peers*. After a connection establishes, the neighbors exchange full BGP routing tables with incremental updates afterward. Neighbors also exchange the KEEPALIVE messages to maintain the connection.

You can classify BGP neighbor routers or peers as internal or external. Connect EBGP peers directly, unless you enable EBGP multihop — IBGP peers do not need direct connection. The IP address of an EBGP neighbor is usually the IP address of the interface directly connected to the router. The BGP process first determines if all internal BGP peers are reachable, then it determines which peers outside the AS are reachable.

1 Assign an AS number, and enter ROUTER-BGP mode from CONFIGURATION mode, from 1 to 65535 for 2-byte, 1 to 4294967295 for 4-byte. Only one AS number is supported per system. If you enter a 4-byte AS number, 4-byte AS support is enabled automatically.
```
router bgp as-number
```
2 Enter a neighbor in ROUTER-BGP mode.
```
neighbor ip-address
```
3 Add a remote AS in ROUTER-NEIGHBOR mode, from 1 to 65535 for 2-byte or 1 to 4294967295 for 4-byte.
```
remote-as as-number
```
4 Enable the BGP neighbor in ROUTER-NEIGHBOR mode.
```
no shutdown
```
5 (Optional) Add a description text for the neighbor in ROUTER-NEIGHBOR mode.
```
description text
```

To reset the configuration when you change the configuration of a BGP neighbor, use the `clear ip bgp *` command. To view the BGP status, use the `show ip bgp summary` command.

Configure BGP

```
OS10# configure terminal
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# neighbor 5.1.1.1
OS10(config-router-neighbor)# remote-as 1
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# description  n1_abcd
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# template t1
OS10(config-router-template)# description peer_template_1_abcd
```

View BGP summary with 2-byte AS number

```
OS10# show ip bgp summary

BGP router identifier 202.236.164.86 local AS number 64901
Neighbor AS MsgRcvd MsgSent Up/Down State/Pfx
120.10.1.1 64701 664 662 04:47:52 established 12000
```

View BGP summary with 4-byte AS number

```
OS10# show ip bgp summary
BGP router identifier 11.1.1.1, local AS number 4294967295
BGP local RIB : Routes to be Added 0, Replaced 0, Withdrawn 0
1 neighbor(s) using 8192 bytes of memory

Neighbor AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/Pfx
5.1.1.2 4294967295 0 0 0 0 0 00:00:00 Active
```

For the router ID, the system selects the first configured IP address or a random number. To view the status of BGP neighbors, use the `show ip bgp neighbors` command. For BGP neighbor configuration information, use the `show running-config bgp` command.

The example shows two neighbors — one is an external BGP neighbor; and the other is an internal BGP neighbor. The first line of the output for each neighbor displays the AS number and states if the link is external or internal.

The third line of the `show ip bgp neighbors` output contains the BGP state. If anything other than *established* displays, the neighbor is not exchanging information and routes. For more information, see IPv6 commands.

View BGP neighbors

```
OS10# show ip bgp neighbors
BGP neighbor is 5.1.1.1, remote AS 1, internal link
BGP version 4, remote router ID 6.1.1.1
BGP state established, in this state for 00:03:11
Last read 01:08:40 seconds, hold time is 180, keepalive interval is 60 seconds
Received 11 messages
3 opens, 1 notifications, 3 updates
4 keepalives, 0 route refresh requests
Sent 14 messages
3 opens, 1 notifications, 0 updates
10 keepalives, 0 route refresh requests

Minimum time between advertisement runs is 0 seconds
Description: n1_abcd
Capabilities received from neighbor for IPv4 Unicast:
MULTIPROTO_EXT(1)ROUTE_REFRESH(2)CISCO_ROUTE_REFRESH(128)
Capabilities advertised to neighbor for IPv4 Unicast:
MULTIPROTO_EXT(1)ROUTE_REFRESH(2)CISCO_ROUTE_REFRESH(128)

Prefixes accepted 3, Prefixes advertised 0

Connections established 3; dropped 2
Closed by neighbor sent 00:03:26 ago

Local host: 5.1.1.2, Local port: 43115
Foreign host: 5.1.1.1, Foreign port: 179
```

View BGP running configuration

```
OS10# show running-configuration bgp
!
router bgp 100
 !
 neighbor 5.1.1.1
 description n1_abcd
```

Configuring BGP in a non-default VRF instance

To configure BGP in a non-default VRF instance.

1 Assign an AS number, and enter ROUTER-BGP mode from CONFIGURATION mode (1 to 65535 for 2-byte, 1 to 4294967295 for 4-byte). Only one AS number is supported per system. If you enter a 4-byte AS number, 4-byte AS support is enabled automatically.

```
router bgp as-number
```

2 Enter ROUTER-BGP-VRF mode to configure BGP in a non-default VRF instance.

```
vrf vrf-name
```

3 Enter a neighbor in CONFIG-ROUTER-VRF mode.

```
neighbor ip-address
```

4 Add a remote AS in ROUTER-NEIGHBOR mode, from 1 to 65535 for 2-byte or 1 to 4294967295 for 4-byte.

```
remote-as as-number
```

5 Enable the BGP neighbor in ROUTER-NEIGHBOR mode.

```
no shutdown
```

6 (Optional) Add a description text for the neighbor in ROUTER-NEIGHBOR mode.

```
description text
```

To reset the configuration when you change the configuration of a BGP neighbor, use the `clear ip bgp *` command. To view the BGP status, use the `show ip bgp summary` command.

Configure BGP

```
OS10# configure terminal
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# vrf blue
OS10(config-router-vrf)# neighbor 5.1.1.1
OS10(config-router-neighbor)# remote-as 1
OS10(config-router-neighbor)# description  n1_abcd
OS10(config-router-neighbor)# exit
OS10(config-router-vrf)# template t1
OS10(config-router-template)# description peer_template_1_abcd
```

Configure Dual Stack

OS10 supports dual stack for BGPv4 and BGPv6. Dual stack BGP allows simultaneous exchange of same IPv4 or IPv6 prefixes through different IPv4 and IPv6 peers. You can enable dual stack using the `activate` command in the corresponding address-family mode. By default, `activate` command is enabled for the IPv4 address family for all the neighbors.

If a BGP-v4 neighbor wants to carry ipv6 prefix information, it activates the IPv6 address-family. For a BGP-v6 neighbor to carry ipv4 prefix, it activates the IPv4 address-family.

1 Enable support for the IPv6 unicast family in CONFIG-ROUTER-BGP mode.

```
address family ipv6 unicast
```

2 Enable IPv6 unicast support on a BGP neighbor/template in CONFIG-ROUTER-BGP-AF mode.

```
activate
```

Configure administrative distance

Routers use administrative distance to determine the best path between two or more routes to reach the same destination. Administrative distance indicates the reliability of the route; the lower the administrative distance, the more reliable the route. If the routing table manager (RTM) receives route updates from one or more routing protocols for a single destination, it chooses the best route based on the administrative distance.

You can assign an administrative distance for the following BGP routes using the `distance bgp` command:

- External BGP (eBGP) routes
- Internal BGP (iBGP) routes
- Local routes

If you do not configure the administrative distance for BGP routes, the following default values are used:

- eBGP—20
- iBGP—200
- local routes—200

To change the administrative distance for BGP, use the following command:

`distance bgp` *`external-distance internal-distance local-distance`*

Configure administrative distance

1 Enable BGP and assign the AS number in CONFIGURATION mode, from 0.1 to 65535.65535 or 1 to 4294967295.

```
OS10# configure terminal
OS10(config)# router bgp 100
```

2 Enter ADDRESS-FAMILY mode.

IPv4:

```
OS10(config-router-bgp-100)# address-family ipv4 unicast
OS10(configure-router-bgpv4-af)#
```

IPv6:

```
OS10(config-router-bgp-100)# address-family ipv6 unicast
OS10(configure-router-bgpv6-af)#
```

3 Change the administrative distance for BGP.

IPv4:

```
OS10(configure-router-bgpv4-af)# distance bgp 21 200 200
```

IPv6:

```
OS10(configure-router-bgpv6-af)# distance bgp 21 201 250
```

The example below provides the configuration for non-default VRF.

```
OS10(config-router-bgp-100)# vrf blue
OS10(config-router-bgp-100-vrf)# address-family ipv4 unicast
OS10(configure-router-bgpv4-vrf-af)# distance bgp 21 200 200
OS10(config-router-bgp-100-vrf)# address-family ipv6 unicast
OS10(configure-router-bgpv6-vrf-af)# distance bgp 21 201 250
```

Peer templates

To configure multiple BGP neighbors at one time, create and populate a BGP peer template. An advantage of configuring peer templates is that members of a peer template inherit the configuration properties of the template and share update policy. Always create a peer template and assign a name to it before adding members to the peer template. Create a peer template before configuring any route policies for the template.

1 Enable BGP and assign the AS number to the local BGP speaker in CONFIGURATION mode, from 1 to 65535 for 2 byte, 1 to 4294967295 | 0.1 to 65535.65535 for 4 byte, or 0.1 to 65535.65535 in dotted format.

```
router bgp as-number
```

2 Create a peer template by assigning a neighborhood name to it in ROUTER-BGP mode.

```
template template-name
```

3 (Optional) Add a text description for the template in ROUTER-TEMPLATE mode.

```
description text
```

4 Enter Address Family mode in ROUTER-NEIGHBOR mode.

```
address-family {[ipv4 | ipv6] [unicast]}
```

5 Filter networks in routing updates, create a route-map and assign a filtering criteria in ROUTER-BGP-NEIGHBOR-AF mode.

```
distribute-list prefix-list-name {in | out}
```

```
route-map map-name {in | out}
```

6 Add a neighbor as a remote AS in ROUTER-TEMPLATE mode, from 1 to 65535 for 2 byte, 1 to 4294967295 | 0.1 to 65535.65535 for 4 byte, or 0.1 to 65535.65535 in dotted format.

```
neighbor ip-address
```

7 Add a remote neighbor, and enter the AS number in ROUTER-TEMPLATE mode.

```
remote-as as-number
```

- To add an EBGP neighbor, configure the `as-number` parameter with a number different from the BGP as-number configured in the `router bgp as-number` command.
- To add an IBGP neighbor, configure the `as-number` parameter with the same BGP as-number configured in the `router bgp as-number` command.

8 Assign a peer-template with a peer-group name from which to inherit to the neighbor in ROUTER-NEIGHBOR mode.

```
inherit template template-name
```

9 Enable the neighbor in ROUTER-BGP mode.

```
no shutdown
```

When you add a peer to a peer group, it inherits all the peer group configured parameters. When you disable a peer group, all the peers within the peer template that are in the `Established` state move to the `Idle` state. A neighbor cannot become a part of a peer group if it has any of these commands configured:

- `advertisement-interval`
- `next-hop-self`
- `route-map out`
- `route-reflector-client`
- `send-community`

A neighbor may keep its configuration after it is added to a peer group if the neighbor configuration is more specific than the peer group and if the neighbor configuration does not affect outgoing updates.

To display the peer-group configuration assigned to a BGP neighbor, enter the `show ip bgp peer-group peer-group-name` command. The `show ip bgp neighbor` command output does not display peer-group configurations.

The following example shows a sample configuration:

Configure peer templates

```
OS10# configure terminal
OS10(config)# router bgp 64601
OS10(config-router-bgp-64601)# template leaf_v4_ebgp
OS10(config-router-template)# description peer_template_1_abcd
OS10(config-router-template)# address-family ipv4 unicast
OS10(config-router-bgp-template-af)# distribute-list leaf_v4_in in
OS10(config-router-bgp-template-af)# distribute-list leaf_v4_out out
OS10(config-router-bgp-template-af)# route-map set_aspath_prepend in
OS10(config-router-bgp-template-af)# exit
OS10(config-router-template)# exit
OS10(config-router-bgp-64601)# neighbor 100.5.1.1
OS10(config-router-neighbor)# inherit template leaf_v4
OS10(config-router-neighbor)# remote-as 64802
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-64601)# neighbor 100.6.1.1
OS10(config-router-neighbor)# inherit template leaf_v4
OS10(config-router-neighbor)# remote-as 64802
OS10(config-router-neighbor)# no shutdown
```

View peer group status

```
OS10# show ip bgp peer-group leaf_v4
Peer-group leaf_v4, remote AS 0
  BGP version 4
  Minimum time between advertisement runs is 30 seconds
  Description: peer_template_1_abcd
  For address family:  Unicast
  BGP neighbor is leaf_v4, peer-group external
  Update packing has 4_OCTET_AS support enabled

  Number of peers in this group 2
  Peer-group members:
    100.5.1.1
    100.6.1.1
```

```
OS10# show ip bgp peer-group leaf_v4 summary
BGP router identifier 100.0.0.8 local AS number 64601
Neighbor        AS       MsgRcvd     MsgSent     Up/Down        State/Pfx
```

```
100.5.1.1       64802    376         325         04:28:25       1251
100.6.1.1       64802    376         327         04:26:17       1251
```

View running configuration

```
OS10# show running-configuration bgp
!
router bgp 64601
 bestpath as-path multipath-relax
 bestpath med missing-as-worst
 non-deterministic-med
 router-id 100.0.0.8
 !
 template leaf_v4
 description peer_template_1_abcd  !
  address-family ipv4 unicast
   distribute-list leaf_v4_in in
   distribute-list leaf_v4_out out
   route-map set_aspath_prepend in
 !
 neighbor 100.5.1.1
  description leaf_connected_ebgp_neighbor
  bfd
  inherit template leaf_v4
  remote-as 64802
  no shutdown
 !
 neighbor 100.6.1.1
  description leaf_connected_ebgp_neighbor
  bfd
  inherit template leaf_v4
  remote-as 64802
  no shutdown
 !
```

Peer templates for a non-default VRF instance

You can create peer templates to add multiple neighbors at a time to the non-default VRF instance that you create.

1 Enable BGP, and assign the AS number to the local BGP speaker in CONFIGURATION mode, from 1 to 65535 for 2 byte, 1 to 4294967295 | 0.1 to 65535.65535 for 4 byte, or 0.1 to 65535.65535 in dotted format.

```
router bgp as-number
```

2 Enter the CONFIG-ROUTER-VRF mode to create a peer template for the non-default VRF instance that you create.

```
vrf vrf-name
```

3 Create a peer template by assigning a neighborhood name to it in CONFIG-ROUTER-VRF mode.

```
template template-name
```

4 Add a neighbor as a remote AS in ROUTER-TEMPLATE mode, from 1 to 65535 for 2 byte, 1 to 4294967295 | 0.1 to 65535.65535 for 4 byte, or 0.1 to 65535.65535 in dotted format.

```
neighbor ip-address
```

5 Add a remote neighbor, and enter the AS number in ROUTER-TEMPLATE mode.

```
remote-as as-number
```

- To add an EBGP neighbor, configure the `as-number` parameter with a number different from the BGP as-number configured in the `router bgp as-number` command.
- To add an IBGP neighbor, configure the `as-number` parameter with the same BGP as-number configured in the `router bgp as-number` command.

6 (Optional) Add a text description for the template in ROUTER-TEMPLATE mode.

```
description text
```

7 Assign a peer-template with a peer-group name from which to inherit to the neighbor in ROUTER-NEIGHBOR mode.

```
inherit template template-name
```

8 Enable the neighbor in ROUTER-BGP mode.

```
neighbor ip-address
```

9 Enable the peer-group in ROUTER-NEIGHBOR mode.

```
no shutdown
```

When you add a peer to a peer group, it inherits all the peer group configured parameters. When you disable a peer group, all the peers within the peer template that are in the Established state move to the Idle state. A neighbor cannot become a part of a peer group if it has any of these commands configured:

- `advertisement-interval`
- `next-hop-self`
- `route-map out`
- `route-reflector-client`
- `send-community`

A neighbor may keep its configuration after it is added to a peer group if the neighbor configuration is more specific than the peer group and if the neighbor configuration does not affect outgoing updates.

To display the peer-group configuration assigned to a BGP neighbor, enter the `show ip bgp peer-group peer-group-name` command. The `show ip bgp neighbor` command output does not display peer-group configurations.

Configure peer templates

```
OS10(config)# router bgp 300
OS10(config-router-bgp-300) vrf blue
OS10(config-router-vrf)# template ebgppg
OS10(config-router-template)# remote-as 100
OS10(config-router-template)# description peer_template_1_abcd
OS10(config-router-template)# exit
OS10(config-router-vrf)# neighbor 3.1.1.1
OS10(config-router-neighbor)# inherit template ebgppg
OS10(config-router-neighbor)# no shutdown
```

Neighbor fall-over

The BGP neighbor fall-over feature reduces the convergence time while maintaining stability. When you enable fall-over, BGP tracks IP reachability to the peer remote address and the peer local address.

When remote or peer local addresses become unreachable, BGP brings the session down with the peer. For example, if no active route exists in the routing table for peer IPv6 destinations/local address, BGP brings the session down.

By default, the hold time governs a BGP session. Configure BGP fast fall-over on a per-neighbor or peer-group basis. BGP routers typically carry large routing tables as frequent session resets are not desirable. If you enable fail-over, the connection to an internal BGP peer is immediately reset if the host route added to reach the internal peer fails.

1 Enter the neighbor IP address in ROUTER-BGP mode.

```
neighbor ip-address
```

2 Disable fast fall-over in ROUTER-NEIGHBOR mode.

```
no fall-over
```

3 Enter the neighbor IP address in ROUTER-BGP mode.

```
neighbor ip-address
```

4 Enable BGP fast fall-Over in ROUTER-NEIGHBOR mode.

```
fall-over
```

Configure neighbor fall-over

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# neighbor 3.1.1.1
```

```
OS10(config-router-neighbor)# remote-as 100
OS10(config-router-neighbor)# fall-over
OS10(config-router-neighbor)# no shutdown
```

Verify neighbor fall-over on neighbor

```
OS10(config-router-neighbor)# do show ip bgp neighbors 3.1.1.1
BGP neighbor is 3.1.1.1, remote AS 100, local AS 100  internal link

  BGP version 4, remote router ID 3.3.3.33
  BGP state ESTABLISHED, in this state for 00:17:17
  Last read 00:27:54 seconds
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time is 180, keepalive interval is 60 seconds
  Fall-over enabled

  Received 23 messages
    1 opens, 0 notifications, 1 updates
    21 keepalives, 0 route refresh requests
  Sent 21 messages
    1 opens, 0 notifications, 0 updates
    20 keepalives, 0 route refresh requests
  Minimum time between advertisement runs is 30 seconds
  Minimum time before advertisements start is 0 seconds
  Capabilities received from neighbor for IPv4 Unicast:
    MULTIPROTO_EXT(1)
    ROUTE_REFRESH(2)
    CISCO_ROUTE_REFRESH(128)
    4_OCTET_AS(65)
  Capabilities advertised to neighbor for IPv4 Unicast:
    MULTIPROTO_EXT(1)
    ROUTE_REFRESH(2)
    CISCO_ROUTE_REFRESH(128)
    4_OCTET_AS(65)
  Prefixes accepted 3, Prefixes advertised 0
  Connections established 1; dropped 0
  Last reset never
  For address family: IPv4 Unicast
    Allow local AS number 0 times in AS-PATH attribute
   Prefixes ignored due to:
    Martian address  0, Our own AS in AS-PATH 0
    Invalid Nexthop  0, Invalid AS-PATH length 0
    Wellknown community 0, Locally originated 0

  For address family: IPv6 Unicast
    Allow local AS number 0 times in AS-PATH attribute
  Local host: 3.1.1.3, Local port: 58633
  Foreign host: 3.1.1.1, Foreign port: 179
```

Verify neighbor fall-over on peer-group

```
OS10# show running-configuration

 !
router bgp 102
 !
 address-family ipv4 unicast
  aggregate-address 6.1.0.0/16
 !
 neighbor 40.1.1.2
  inherit template bgppg
  no shutdown
 !
 neighbor 60.1.1.2
  inherit template bgppg
  no shutdown
 !
 neighbor 32.1.1.2
  remote-as 100
```

```
  no shutdown
 !
 template bgppg
  fall-over
  remote-as 102
!
```

Configure password

You can enable message digest 5 (MD5) authentication with a password on the TCP connection between two BGP neighbors.

Configure the same password on both BGP peers. When you configure MD5 authentication between two BGP peers, each segment of the TCP connection is verified and the MD5 digest is checked on every segment sent on the TCP connection. Configuring a password for a neighbor establishes a new connection.

NOTE: You can secure the VTEP neighbor communications as well using the MD5 authentication.

Configure password

- Configure the password in both the BGP peers in ROUTER-NEIGHBOR CONFIGURATION or ROUTER-TEMPLATE CONFIGURATION mode. The password provided in ROUTER-NEIGHBOR mode takes preference over the password in ROUTER-TEMPLATE mode. Enter the password either as plain text or in encrypted format.
  - `password {9 encrypted password-string|password-string}`

View password configuration

- `show configuration`

Peer 1 in ROUTER-NEIGHBOR mode

```
OS10# configure terminal
OS10(config)# interface ethernet 1/1/5
OS10(conf-if-eth1/1/5)# no switchport
OS10(conf-if-eth1/1/5)# ip address 11.1.1.1/24
OS10(conf-if-eth1/1/5)# router bgp 10
OS10(config-router-bgp-10)# neighbor 11.1.1.2
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# remote-as 10
OS10(config-router-neighbor)# password abcdell
```

Peer 1 in ROUTER-TEMPLATE mode

```
OS10# configure terminal
OS10(config)# interface ethernet 1/1/5
OS10(conf-if-eth1/1/5)# no switchport
OS10(conf-if-eth1/1/5)# ip address 11.1.1.1/24
OS10(conf-if-eth1/1/5)# router bgp 10
OS10(config-router-bgp-10)# template pass
OS10(config-router-template)# password 9
f785498c228f365898c0efdc2f476b4b27c47d972c3cd8cd9b91f518c14ee42d
OS10(config-router-template)# exit
OS10(config-router-bgp-10)# neighbor 11.1.1.2
OS10(config-router-neighbor)# inherit template pass
```

View password configuration in peer 1

```
OS10(config-router-neighbor)# show configuration
!
neighbor 11.1.1.2
 password 9 0fbe1ad397712f74f4df903b4ff4b7b6e22cc377180432d7523a70d403d41565
 remote-as 10
 no shutdown
```

```
OS10(config-router-neighbor)# do show running-configuration bgp
!
```

```
router bgp 10
 !
 template pass
  password 9 f785498c228f365898c0efdc2f476b4b27c47d972c3cd8cd9b91f518c14ee42d
 !
 neighbor 11.1.1.2
  inherit template pass
  password 9 01320afb39f49134882b0a9814fe6e8e228f616f60a35958844775314c00f0e5
  remote-as 10
  no shutdown
```

Peer 2 in ROUTER-NEIGHBOR mode

```
OS10# configure terminal
OS10(config)# interface ethernet 1/1/5
OS10(conf-if-eth1/1/5)# no switchport
ip OS10(conf-if-eth1/1/5)# ip address 11.1.1.2/24
OS10(conf-if-eth1/1/5)# router bgp 20
OS10(config-router-bgp-20)# neighbor 11.1.1.1
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# remote-as 20
OS10(config-router-neighbor)# password abcdell
```

Peer 2 in ROUTER-TEMPLATE mode

```
OS10# configure terminal
OS10(config)# interface ethernet 1/1/5
OS10(conf-if-eth1/1/5)# no switchport
OS10(conf-if-eth1/1/5)# ip address 11.1.1.2/24
OS10(conf-if-eth1/1/5)# router bgp 20
OS10(config-router-bgp-20)# template pass
OS10(config-router-template)# password 9
f785498c228f365898c0efdc2f476b4b27c47d972c3cd8cd9b91f518c14ee42d
OS10(config-router-template)# exit
OS10(config-router-bgp-20)# neighbor 11.1.1.1
OS10(config-router-neighbor)# inherit template pass
```

View password configuration in peer 2

```
OS10(config-router-neighbor)# show configuration
!
neighbor 11.1.1.1
 password 9 0fbe1ad397712f74f4df903b4ff4b7b6e22cc377180432d7523a70d403d41565
 remote-as 20
 no shutdown
```

```
OS10(config-router-neighbor)# do show running-configuration bgp
!
router bgp 20
  neighbor 11.1.1.2
  password 9 f785498c228f365898c0efdc2f476b4b27c47d972c3cd8cd9b91f518c14ee42d
  remote-as 20
  no shutdown
```

Fast external fallover

Fast external fallover terminates EBGP sessions of any directly adjacent peer if the link used to reach the peer goes down. BGP does not wait for the hold-down timer to expire.

Fast external fallover is enabled by default. To disable or re-enable it, use the `[no] fast-external-fallover` command. For the `fast-external-fallover` command to take effect on an established BGP session, you must reset the session using the `clear ip bgp {* | peer-ipv4-address | peer-ipv6-address}` command.

View fast external fallover configuration

```
OS10(config)# do show running-configuration bgp
!
router bgp 300
 !
 neighbor 3.1.1.1
  remote-as 100
  no shutdown
 !
 neighbor 3::1
  remote-as 100
  no shutdown
  !
  address-family ipv6 unicast
   activate
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# show configuration
!
interface ethernet1/1/1
 ip address 3.1.1.3/24
 no switchport
 no shutdown
 ipv6 address 3::3/64
OS10(conf-if-eth1/1/1)# shutdown
OS10(conf-if-eth1/1/1)# do show ip bgp summary
BGP router identifier 11.11.11.11 local AS number 300
Neighbor                                        AS          MsgRcvd     MsgSent     Up/
Down           State/Pfx
3.1.1.1                                         100         6           6
00:00:15          Active
3::1                                            100         8           11
00:00:15          Active
OS10(conf-if-eth1/1/1)#
```

View fast external fallover unconfiguration

```
OS10(config-router-bgp-300)# do show running-configuration bgp
!
router bgp 300
 no fast-external-fallover
 !
 neighbor 3.1.1.1
  remote-as 100
  no shutdown
 !
 neighbor 3::1
  remote-as 100
  no shutdown
  !
  address-family ipv6 unicast
   activate
OS10(config-router-bgp-300)#
OS10(conf-if-eth1/1/1)# do clear ip bgp *
OS10# show ip bgp summary
BGP router identifier 11.11.11.11 local AS number 300
Neighbor  AS          MsgRcvd     MsgSent     Up/Down     State/Pfx
------------------------------------------------------------------
3.1.1.1   100         7           4           00:00:08    3
3::1      100         9           5           00:00:08    4
OS10#
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# shutdown
OS10(conf-if-eth1/1/1)# do show ip bgp summary
BGP router identifier 11.11.11.11 local AS number 300
Neighbor  AS          MsgRcvd     MsgSent     Up/Down     State/Pfx
------------------------------------------------------------------
3.1.1.1   100         7           4           00:00:29    3
3::1      100         9           5           00:00:29    4
```

```
OS10(conf-if-eth1/1/1)#
OS10(config-router-bgp-neighbor-af)# Apr 27 01:39:03 OS10 dn_sm[2065]: Node.1-Unit.1:PRI:alert
[os10:event],
 %Dell EMC (OS10) %BGP_NBR_BKWD_STATE_CHG: Backward state change occurred Hold Time expired for
Nbr:3.1.1.3 VRF:default
 Apr 27 01:39:03 OS10 dn_sm[2065]: Node.1-Unit.1:PRI:alert [os10:event], %Dell EMC (OS10)
%BGP_NBR_BKWD_STATE_CHG: Backward
 state change occurred Hold Time expired for Nbr:3::3 VRF:default
```

Passive peering

When you enable a peer-template, the system sends an OPEN message to initiate a TCP connection. If you enable passive peering for the peer template, the system does not send an OPEN message but responds to an OPEN message.

When a BGP neighbor connection with authentication rejects a passive peer-template, the system prevents another passive peer-template on the same subnet from connecting with the BGP neighbor. To work around this constraint, change the BGP configuration or change the order of the peer template configuration.

You can restrict the number of passive sessions the neighbor accepts using the `limit` command.

1 Enable BGP and assign the AS number to the local BGP speaker in CONFIGURATION mode (1 to 65535 for 2-byte, 1 to 4294967295 for 4-byte).

```
router bgp as-number
```

2 Configure a template that does not initiate TCP connections with other peers in ROUTER-BGP mode. A maximum of 16 characters.

```
template template-name
```

3 Create and enter the AS number for the remote neighbor in ROUTER-BGP-TEMPLATE mode (1 to 4294967295).

```
remote-as as-number
```

4 Enable peer listening and enter the maximum dynamic peers count in ROUTER-BGP-TEMPLATE mode (1 to 4294967295).

```
listen neighbor ip-address limit
```

Only after the peer template responds to an OPEN message sent on the subnet does the state of its BGP change to ESTABLISHED. After the peer template is ESTABLISHED, the peer template is the same as any other peer template, see Peer templates.

If you do not configure a BGP device in Peer-Listening mode, a session with a dynamic peer comes up. Passwords are not supported on BGPv4/v6 dynamic peers.

Configure passive peering

```
OS10(config)# router bgp 10
OS10(conf-router-bgp-10)# template bgppg
OS10(conf-router-template)# remote-as 100
OS10(conf-router-template)# listen 32.1.0.0/8 limit 10
```

Local AS

During BGP network migration, you can maintain existing AS numbers. Reconfigure your routers with the new information to disable after the migration. Network migration is not supported on passive peer templates. You must configure Peer templates before assigning it to an AS.

1 Enter a neighbor IP address, A.B.C.D, in ROUTER-BGP mode.

```
neighbor ip-address
```

2 Enter a local-as number for the peer, and the AS values not prepended to announcements from the neighbors in ROUTER-NEIGHBOR mode (1 to 4294967295).

```
local-as as number [no prepend]
```

3 Return to ROUTER-BGP mode.

```
exit
```

4 Enter a template name to assign to the peer-groups in ROUTER-BGP mode. A maximum of 16 characters.

```
template template-name
```

5 Enter a local-as number for the peer in ROUTER-TEMPLATE mode.

```
local-as as number [no prepend]
```

6 Add a remote AS in ROUTER-TEMPLATE mode (1 to 65535 for 2 bytes, 1 to 4294967295 for 4 bytes).

```
remote-as as-number
```

Allow external routes from neighbor

```
OS10(config)# router bgp 10
OS10(conf-router-bgp-10)# neighbor 32.1.1.2
OS10(conf-router-neighbor)# local-as 50
OS10(conf-router-neighbor)# exit
OS10(conf-router-bgp-10)# template bgppg1
OS10(conf-router-template)# fall-over
OS10(conf-router-template)# local-as 400
OS10(conf-router-template)# remote-as 102
```

Local AS number disabled

```
OS10(config)# router bgp 102
OS10(conf-router-bgp-102)# neighbor 32.1.1.2
OS10(conf-router-neighbor)# no local-as 100
```

AS number limit

Sets the number of times an AS number occurs in an AS path. The `allow-as` parameter permits a BGP speaker to allow the AS number for a configured number of times in the updates received from the peer.

The AS-PATH loop is detected if the local AS number is present more than the number of times in the command.

1 Enter the neighbor IP address to use the AS path in ROUTER-BGP mode.

```
neighbor ip address
```

2 Enter Address Family mode in ROUTER-NEIGHBOR mode.

```
address-family {[ipv4 | ipv6] [unicast])
```

3 Allow the neighbor IP address to use the AS path the specified number of times in ROUTER-BGP-NEIGHBOR-AF mode (1 to 10).

```
allowas-in number
```

Configure AS number appearance

```
OS10(config)# router bgp 10
OS10(conf-router-bgp-10)# neighbor 1.1.1.2
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-bgp-neighbor-af)# allowas-in 5
```

View AS numbers in AS paths

```
OS10# show running-configuration bgp
!
router bgp 101
 no fast-external-fallover
 !
 address-family ipv4 unicast
  dampening
 !
 neighbor 17.1.1.2
  remote-as 102
  no shutdown
  !
  address-family ipv4 unicast
   allowas-in 4
```

Show IP BGP

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# neighbor 172:16:1::2
OS10(config-router-neighbor)# remote-as 100
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv6 unicast
OS10(config-router-bgp-neighbor-af)# activate
OS10(config-router-bgp-neighbor-af)# allowas-in 1
OS10(config-router-bgp-neighbor-af)# end
OS10# show running-configuration bgp
!
router bgp 100
 !
 neighbor 172:16:1::2
  remote-as 100
  no shutdown
  !
  address-family ipv6 unicast
   activate
   allowas-in 1
OS10# show ip bgp
BGP local RIB : Routes to be Added , Replaced , Withdrawn
BGP local router ID is 100.1.1.1
Status codes: s suppressed, S stale, d dampened, h history, * valid, > best
Path source: I - internal, a - aggregate, c - confed-external,
r - redistributed/network, S - stale
Origin codes: i - IGP, e - EGP, ? - incomplete
          Network          Next Hop      Metric     LocPrf     Weight     Path
*>I       55::/64          172:16:1::2   0          0          0          100 200 300 400 i
*>I       55:0:0:1::/64    172:16:1::2   0          0          0          100 200 300 400 i
*>I       55:0:0:2::/64    172:16:1::2   0          0          0          100 200 300 400 i
```

Redistribute routes

Add routes from other routing instances or protocols to the BGP process. You can include OSPF, static, or directly connected routes in the BGP process with the `redistribute` command.

- Include directly connected or user-configured (static) routes in ROUTER-BGP-AF mode.
  ```
  redistribute {connected | static}
  ```
- Include specific OSPF routes in IS-IS in ROUTER-BGP-AF mode (1 to 65535).
  ```
  redistribute ospf process-id
  ```

Disable redistributed routes

```
OS10(conf-router-bgp-af)# no redistribute ospf route-map ospf-to-bgp
```

Enable redistributed routes

```
OS10(conf-router-bgp-af)# redistribute ospf
```

Additional paths

The `add-path` command is disabled by default.

1 Assign an AS number in CONFIGURATION mode.
  ```
  router bgp as-number
  ```
2 Enter a neighbor and IP address (A.B.C.D) in ROUTER-BGP mode.
  ```
  neighbor ip-address
  ```

3 Enter Address Family mode in ROUTER-NEIGHBOR mode.

```
address-family {[ipv4 | ipv6] [unicast])
```

4 Allow the specified neighbor to send or receive multiple path advertisements in ROUTER-BGP mode. The *count* parameter controls the number of paths that are advertised — not the number of paths received.

```
add-path [both | received | send] count
```

Enable additional paths

```
OS10(config)# router bgp 102
OS10(conf-router-bgp-102)# neighbor 32.1.1.2
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-bgp-neighbor-af)# add-path both 3
```

MED attributes

OS10 uses the MULTI_EXIT_DISC or MED attribute when comparing EBGP paths from the same AS. MED comparison is not performed in paths from neighbors with different AS numbers.

1 Enable MED comparison in the paths from neighbors with different AS in ROUTER-BGP mode.

```
always-compare-med
```

2 Change the best path MED selection in ROUTER-BGP mode.

```
bestpath med {confed | missing-as-best}
```

- `confed`—Selects the best path MED comparison of paths learned from BGP confederations.
- `missing-as-best`—Treats a path missing an MED as the most preferred one.
- `missing-as-worst`—Treats a path missing an MED as the least preferred one.

Modify MED attributes

```
OS10(config)# router bgp 100
OS10(conf-router-bgp-100)# always-compare-med
OS10(conf-router-bgp-100)# bestpath med confed
```

Local preference attribute

You can change the value of the LOCAL_PREFERENCE attributes for all routes the router receives. To change the LOCAL_PREF value in ROUTER-BGP mode from 0 to 4294967295 with default 100, use the `default local preference value` command.

To view the BGP configuration, use the `show running-configuration` command. A more flexible method for manipulating the LOCAL_PREF attribute value is to use a route-map.

1 Assign a name to a route map in CONFIGURATION mode.

```
route-map map-name {permit | deny | sequence-number}
```

2 Change the LOCAL_PREF value for routes meeting the criteria of this route map in ROUTE-MAP mode, then return to CONFIGURATION mode.

```
set local-preference value
exit
```

3 Enter ROUTER-BGP mode.

```
router bgp as-number
```

4 Enter the neighbor to apply the route map configuration in ROUTER-BGP mode.

```
neighbor {ip-address}
```

5 Apply the route map to the neighbor's incoming or outgoing routes in ROUTER-BGP-NEIGHBOR-AF mode.

```
route-map map-name {in | out)
```

6 Enter the peer group to apply the route map configuration in ROUTER-BGP mode.

```
template template-name
```

7 Apply the route map to the peer group's incoming or outgoing routes in CONFIG-ROUTER-TEMPLATE-AF mode.

```
route-map map-name {in | out}
```

Configure and view local preference attribute

```
OS10(config)# route-map bgproutemap 1
OS10(conf-route-map)# set local-preference 500
OS10(conf-route-map)# exit
OS10(config)# router bgp 10
OS10(conf-router-bgp-10)# neighbor 10.1.1.4
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-bgp-neighbor-af)# route-map bgproutemap in
```

```
OS10 configure terminal
OS10(config)# route-map bgproutemap 1
OS10(conf-route-map)# set local-preference 500
OS10(conf-route-map)# exit
OS10(config)# router bgp 64601
OS10(conf-router-bgp-64601)# template bgppg
OS10(conf-router-template)# address-family ipv4 unicast
OS10(conf-router-bgp-template-af)# route-map bgproutemap in
```

View route-map

```
OS10(conf-route-map)# do show route-map
route-map bgproutemap, permit, sequence 1
  Match clauses:
  Set clauses:
    local-preference 500
    metric 400
    origin incomplete
```

Weight attribute

You can influence the BGP routing based on the weight value. Routes with a higher weight value have preference when multiple routes to the same destination exist.

1 Assign a weight to the neighbor connection in ROUTER-BGP mode.

```
neighbor {ip-address}
```

2 Set a weight value for the route in ROUTER-NEIGHBOR mode (1 to 4294967295, default 0).

```
weight weight
```

3 Return to ROUTER-BGP mode.

```
exit
```

4 Assign a weight value to the peer-group in ROUTER-BGP mode.

```
template template name
```

5 Set a weight value for the route in ROUTER-TEMPLATE mode.

```
weight weight
```

Modify weight attribute

```
OS10(config)# router bgp 10
OS10(config-router-bgp-10)# neighbor 10.1.1.4
OS10(config-router-neighbor)# weight 400
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-10)# template zanzibar
OS10(config-router-template)# weight 200
```

Enable multipath

You can have one path to a destination by default, and enable multipath to allow up to 64 parallel paths to a destination. The `show ip bgp network` command includes multipath information for that network.

- Enable multiple parallel paths in ROUTER-BGP mode.
  ```
  maximum-paths {ebgp | ibgp} number
  ```

Enable multipath

```
OS10(config)# router bgp 10
OS10(conf-router-bgp-10)# maximum-paths ebgp 10
```

Route-map filters

Filtering routes allows you to implement BGP policies. Use route-maps to control which routes the BGP neighbor or peer group accepts and advertises.

1 Enter the neighbor IP address to filter routes in ROUTER-BGP mode.
  ```
  neighbor ipv4-address
  ```
2 Enter Address Family mode in ROUTER-NEIGHBOR mode.
  ```
  address-family {[ipv4 | ipv6] [unicast]}
  ```
3 Create a route-map and assign a filtering criteria in ROUTER-BGP-NEIGHBOR-AF mode, then return to CONFIG-ROUTER-BGP mode.
  ```
  route-map map-name {in | out}
  exit
  ```
  - `in`—Enter a filter for incoming routing updates.
  - `out`—Enter a filter for outgoing routing updates.
4 Enter a peer template name in ROUTER-BGP mode.
  ```
  template template-name
  ```
5 Enter Address Family mode.
  ```
  address-family {[ipv4 | ipv6] [unicast]}
  ```
6 Create a route-map, and assign a filtering criteria in ROUTER-BGP-TEMPLATE-AF mode.
  ```
  route-map map-name {in | out}
  ```

Filter BGP route

```
OS10(config)# router bgp 102
OS10(conf-router-bgp-102)# neighbor 40.1.1.2
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-bgp-neighbor-af)# route-map metro in
OS10(conf-router-bgp-neighbor-af)# exit
OS10(conf-router-bgp-102)# template ebgp
OS10(conf-router-template)# address-family ipv4 unicast
OS10(conf-router-bgp-template-af)# route-map metro in
```

Route reflector clusters

BGP route reflectors are intended for ASs with a large mesh. They reduce the amount of BGP control traffic. With route reflection configured properly, IBGP routers are not fully meshed within a cluster but all receive routing information.

Configure clusters of routers where one router is a concentration router and the others are clients who receive their updates from the concentration router.

1 Assign an ID to a router reflector cluster in ROUTER-BGP mode. You can have multiple clusters in an AS.
```
cluster-id cluster-id
```
2 Assign a neighbor to the router reflector cluster in ROUTER-BGP mode.
```
neighbor {ip-address}
```
3 Configure the neighbor as a route-reflector client in ROUTER-NEIGHBOR mode, then return to ROUTER-BGP mode.
```
route-reflector-client
exit
```
4 Assign a peer group template as part of the route-reflector cluster in ROUTER-BGP mode.
```
template template-name
```
5 Configure the template as the route-reflector client in ROUTER-TEMPLATE mode.
```
route-reflector-client
```

When you enable a route reflector, the system automatically enables route reflection to all clients. To disable route reflection between all clients in this reflector, use the `no bgp client-to-client reflection` command in ROUTER-BGP mode. You must fully mesh all the clients before you disable route reflection.

Configure BGP route reflector

```
OS10(config)# router bgp 102
OS10(conf-router-bgp-102)# cluster-id 4294967295
OS10(conf-router-bgp-102)# neighbor 32.1.1.2
OS10(conf-router-neighbor)# route-reflector-client
OS10(conf-router-neighbor)# exit
OS10(conf-router-bgp-102)# template zanzibar
OS10(conf-router-template)# route-reflector-client
```

Aggregate routes

OS10 provides multiple ways to aggregate routes in the BGP routing table. At least one route of the aggregate must be in the routing table for the configured aggregate route to become active. AS_SET includes AS_PATH and community information from the routes included in the aggregated route.

1 Assign an AS number in CONFIGURATION mode.
```
router bgp as-number
```
2 Enter Address Family mode in ROUTER-BGP mode.
```
address-family {[ipv4 | ipv6] [unicast]}
```
3 Aggregate address in ROUTER-BGPv4-AF mode.
```
aggregate-address ip-address/mask
```

Configure aggregate routes

```
OS10(config)# router bgp 105
OS10(conf-router-bgp-105)# address-family ipv4 unicast
OS10(conf-router-bgpv4-af)# aggregate-address 3.3.0.0/16
```

View running configuration

```
OS10(conf-router-bgpv4-af)# do show running-configuration bgp
! Version
! Last configuration change at Jul  27 06:51:17 2016
!
!
router bgp 105
 !
 address-family ipv4 unicast
  aggregate-address 3.3.0.0/16
```

```
 !

neighbor 32.1.1.2
  remote-as 104
  no shutdown
  !
  address-family ipv4 unicast
```

Confederations

Another way to organize routers within an AS and reduce the mesh for IBGP peers is to configure BGP confederations. As with route reflectors, Dell EMC recommends BGP confederations only for IBGP peering involving many IBGP peering sessions per router.

When you configure BGP confederations, you break the AS into smaller sub-ASs. To devices outside your network, the confederations appear as one AS. Within the confederation sub-AS, the IBGP neighbors are fully meshed and the MED, NEXT_HOP, and LOCAL_PREF attributes maintain between confederations.

1 Enter the confederation ID AS number in ROUTER-BGP mode (1 to 65535 for 2–byte, 1 to 4294967295 for 4–byte).

```
confederation identifier as-number
```

2 Enter which confederation sub-AS are peers in ROUTER-BGP mode, from 1 to 65535 for 2–byte, 1 to 4294967295 for 4–byte. All Confederation routers must be either 4 bytes or 2 bytes. You cannot have a mix of router ASN support.

```
confederation peers as-number [... as-number]
```

Configure BGP confederations

```
OS10(config)# router bgp 65501
OS10(conf-router-bgp-65501)# confederation identifier 100
OS10(conf-router-bgp-65501)# confederation peers 65502 65503 65504
OS10(conf-router-bgp-65501)# neighbor 1.1.1.2
OS10(conf-router-neighbor)# remote-as 65502
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# exit
OS10(conf-router-bgp-65501)# neighbor 2.1.1.2
OS10(conf-router-neighbor)# remote-as 65503
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# exit
OS10(conf-router-bgp-65501)# neighbor 3.1.1.2
OS10(conf-router-neighbor)# remote-as 65504
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# exit
OS10(conf-router-bgp-65501)# end
OS10# show running-configuration bgp
!
router bgp 65501
 confederation identifier 100
 confederation peers 65502 65503 65504
 !
 neighbor 1.1.1.2
  remote-as 65502
  no shutdown
 !
 neighbor 2.1.1.2
  remote-as 65503
  no shutdown
 !
 neighbor 3.1.1.2
  remote-as 65504
  no shutdown
```

Route dampening

When EBGP routes become unavailable, they "flap" and the router issues both WITHDRAWN and UPDATE notices. A flap occurs when a route is withdrawn, readvertised after being withdrawn, or has an attribute change.

The constant router reaction to the WITHDRAWN and UPDATE notices causes instability in the BGP process. To minimize this instability, configure penalties (a numeric value) for routes that flap. When that penalty value reaches a configured limit, the route is not advertised, even if the route is up, the penalty value is 1024.

As time passes and the route does not flap, the penalty value decrements or decays. If the route flaps again, it is assigned another penalty. The penalty value is cumulative and adds underwithdraw, readvertise, or attribute change.

When dampening applies to a route, its path is described by:

History entry	Entry that stores information on a downed route.
Dampened path	Path that is no longer advertised.
Penalized path	Path that is assigned a penalty.

1 Enable route dampening in ROUTER-BGP mode.

```
dampening [half-life | reuse | max-suppress-time]
```

- *half-life* — Number of minutes after which the penalty decreases (1 to 45, default 15). After the router assigns a penalty of 1024 to a route, the penalty decreases by half after the half-life period expires.
- *reuse* — Number compares to the flapping route's penalty value. If the penalty value is less than the reuse value, the flapping route again advertises or is no longer suppressed (1 to 20000, default 750). Withdrawn routes are removed from the history state.
- *suppress* — Number compares to the flapping route's penalty value. If the penalty value is greater than the suppress value, the flapping route no longer advertises and is suppressed (1 to 20000, default 2000).
- *max-suppress-time* — Maximum number of minutes a route is suppressed (1 to 255, default is four times the half-life value or 60 minutes).

2 View all flap statistics or for specific routes meeting the criteria in EXEC mode.

```
show ip bgp flap-statistics [ip-address [mask]]
```

- *ip-address [mask]* — Enter the IP address and mask.
- `filter-list` *as-path-name* — Enter the name of an AS-PATH ACL.
- `regexp` *regular-expression* — Enter a regular express to match on.

When you change the best path selection method, path selections for the existing paths remain unchanged until you reset it by using the `clear ip bgp` command in EXEC mode.

Configure values to reuse or restart route

```
OS10(config)# router bgp 102
OS10(conf-router-bgp-102)# address-family ipv4 unicast
OS10(conf-router-bgpv4-af)# dampening 2 2000 3000 10
```

View dampened (nonactive) routes

```
OS10# show ip bgp flap-statistics

BGP local router ID is 13.176.123.28
Status codes: s suppressed, S stale, d dampened, h history, * valid, > best
Origin codes: i - IGP, e - EGP, ? - incomplete
          Network
From                                    Reuse          Path
Total number of prefixes: 0
```

View dampened paths

```
OS10# show ip bgp dampened-paths

BGP local router ID is 80.1.1.1
Status codes: s suppressed, S stale, d dampened, h history, * valid, > best
Origin codes: i - IGP, e - EGP, ? - incomplete
      Network         From       Reuse       Path
d*    3.1.2.0/24      80.1.1.2   00:00:12    800 9 8 i
d*    3.1.3.0/24      80.1.1.2   00:00:12    800 9 8 i
d*    3.1.4.0/24      80.1.1.2   00:00:12    800 9 8 i
d*    3.1.5.0/24      80.1.1.2   00:00:12    800 9 8 i
d*    3.1.6.0/24      80.1.1.2   00:00:12    800 9 8 i
Total number of prefixes: 5
```

Timers

To adjust the routing timers for all neighbors, configure the timer values using the `timers` command. If both the peers negotiate with different keepalive and hold time values, the final hold time value is the lowest values received. The new keepalive value is one-third of the accepted hold time value.

- Configure timer values for all neighbors in ROUTER-NEIGHBOR mode.

  ```
  timers keepalive holdtime
  ```

  - *keepalive* — Time interval in seconds, between keepalive messages sent to the neighbor routers (1 to 65535, default 60).
  - *holdtime* — Time interval in seconds, between the last keepalive message and declaring the router dead (3 to 65535, default 180).

View nondefault values

```
OS10# show running-configuration
...
 neighbor 32.1.1.2
 remote-as 103
  timers 61 181
  no shutdown
```

Neighbor soft-reconfiguration

BGP soft-reconfiguration allows for fast and easy route changes. Changing routing policies requires a reset of BGP sessions or the TCP connection, for the policies to take effect.

Resets cause undue interruption to traffic due to the hard reset of the BGP cache, and the time it takes to re-establish the session. BGP soft-reconfiguration allows for policies to apply to a session without clearing the BGP session. You can perform a soft-reconfiguration on a per-neighbor basis, either inbound or outbound. BGP soft-reconfiguration clears the policies without resetting the TCP connection. After configuring soft-reconfiguration, use the `clear ip bgp` command to make the neighbor use soft reconfiguration.

When you enable soft-reconfiguration for a neighbor and you execute the `clear ip bgp soft in` command, the update database stored in the router replays and updates are re-evaluated. With this command, the replay and update process triggers only if a route-refresh request is not negotiated with the peer. If the request is negotiated after using the `clear ip bgp soft in` command, BGP sends a route-refresh request to the neighbor and receives all the peer's updates.

To use soft reconfiguration, or soft reset without preconfiguration, both BGP peers must support soft route refresh. The soft route refresh advertises in the OPEN message sent when the peers establish a TCP session. To determine whether a BGP router supports this capability, use the `show ip bgp neighbors` command. If a router supports the route refresh capability, the `Received route refresh capability from peer` message displays.

1 Enable soft-reconfiguration for the BGP neighbor and BGP template in ROUTER-BGP mode. BGP stores all the updates that the neighbor receives but does not reset the peer-session. Entering this command starts the storage of updates, which is required to do inbound soft reconfiguration.

```
neighbor {ip-address} soft-reconfiguration inbound
```

2 Enter Address Family mode in ROUTER-NEIGHBOR mode.

```
address-family {[ipv4 | ipv6] [unicast]}
```

3 Configure soft-configuration for the neighbors belonging to the template.

```
soft-reconfiguration inbound
```

4 Clear all information or only specific details in EXEC mode.

```
clear ip bgp {neighbor-address | * } [soft in]
```

- `*` — Clears all peers.
- `neighbor-address` — Clears the neighbor with this IP address.

Soft-reconfiguration of IPv4 neighbor

```
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-bgp-neighbor-af)# soft-reconfiguration inbound
```

Soft-reconfiguration of IPv6 neighbor

```
OS10(conf-router-neighbor)# address-family ipv6 unicast
OS10(conf-router-bgp-neighbor-af)# soft-reconfiguration inbound
```

BGP commands

activate

Enables the neighbor or peer group to be the current address-family identifier (AFI).

Syntax	`activate`
Parameters	None
Default	Not configured
Command Mode	ROUTER-BGP-NEIGHBOR-AF
Usage Information	This command exchanges IPv4 or IPv6 address family information with an IPv4 or IPv6 neighbor. IPv4 unicast Address family is enabled by default. To activate IPv6 address family for IPv6 neighbor, use the `activate` command. To de-activate IPv4 address family for IPv6 neighbor, use the `no` activate command.
Example	`OS10(conf-router-neighbor)# address-family ipv4 unicast` `OS10(conf-router-bgp-neighbor-af)# activate`
Supported Releases	10.2.0E or later

add-path

Allows the system to advertise multiple paths for the same destination without replacing previous paths with new ones.

Syntax	`add-path {both path count | receive | send path count}`
Parameters	• `both path count` — Enter the number of paths to advertise to the peer, from 2 to 64.

- `receive` — Receive multiple paths from the peer.
- `send path count` — Enter the number of multiple paths to send multiple to the peer, from 2 to 64.

Default	Not configured
Command Mode	ROUTER-BGP-NEIGHBOR-AF
Usage Information	Advertising multiple paths to peers for the same address prefix without replacing the existing path with a new one reduces convergence times. The `no` version of this command disables the multiple path advertisements for the same destination.
Example (IPv4)	`OS10(conf-router-bgp-af)# add-path both 64`
Example (IPv6)	`OS10(conf-router-bgpv6-af)# add-path both 64`
Example (Receive)	`OS10(conf-router-bgpv6-af)# add-path receive`
Supported Releases	10.2.0E or later

address-family

Enters Global Address Family Configuration mode for the IP address family.

Syntax	`address-family {[ipv4 | ipv6] unicast}`
Parameters	• `ipv4 unicast` — Enter an IPv4 unicast address family. • `ipv6 unicast` — Enter an IPv6 unicast address family.
Default	None
Command Mode	ROUTER-BGP
Usage Information	This command applies to all IPv4 or IPv6 peers belonging to the template or neighbors only. The `no` version of this command removes the subsequent address-family configuration.

Example (IPv4 Unicast)	`OS10(config)# router bgp 3`
	`OS10(conf-router-bgp-3)# address-family ipv4 unicast`
	`OS10(conf-router-bgpv4-af)#`

Example (IPv6 Unicast)	`OS10(config)# router bgp 4`
	`OS10(conf-router-bgp-4)# address-family ipv6 unicast`
	`OS10(conf-router-bgpv6-af)#`

Supported Releases	10.3.0E or later

advertisement-interval

Sets the minimum time interval for advertisement between the BGP neighbors or within a BGP peer group.

Syntax	`advertisement-interval seconds`
Parameters	*seconds*—Enter the time interval value in seconds between BGP advertisements, from 1 to 600.
Default	EBGP 30 seconds, IBGP 5 seconds
Command Mode	ROUTER-NEIGHBOR

Usage Information	The time interval applies to all peer group members of the template in ROUTER-TEMPLATE mode. The `no` version of this command resets the advertisement-interval value to the default.
Example	`OS10(conf-router-neighbor)# advertisement-interval 50`
Supported Releases	10.3.0E or later

advertisement-start

Delays initiating the OPEN message for the specified time.

Syntax	`advertisement-start` *`seconds`*
Parameters	*`seconds`*—Enter the time interval value, in seconds, before starting to send the BGP OPEN message, from 0 to 240.
Default	Not configured
Command Mode	ROUTER-NEIGHBOR
Usage Information	The time interval applies to all the peer group members of the template in ROUTER-TEMPLATE mode. The `no` version of this command disables the advertisement-start time interval.
Example	`OS10(conf-router-neighbor)# advertisement-start 30`
Supported Releases	10.3.0E or later

aggregate-address

Summarizes a range of prefixes to minimize the number of entries in the routing table.

Syntax	`aggregate-address` *`address/mask`* `[as-set] [summary-only] [advertise-map` *`map-name`*`] {attribute-map` *`route-map-name`*`] [suppress-map` *`route-map-name`*`]`
Parameters	• *`address/mask`* — Enter the IP address and mask. • `as-set` — (Optional) Generates AS set-path information. • `summary-only` — (Optional) Filters more specific routes from updates. • `advertise-map` *`map-name`* — (Optional) Enter the map name to advertise. • `attribute-map` *`route-map-name`* — (Optional) Enter the route-map name to set aggregate attributes. • `suppress-map` *`route-map-name`* — (Optional) Enter the route-map name to conditionally filters specific routes from updates.
Default	None
Command Mode	ROUTER-BGPv4-AF
Usage Information	At least one of the routes included in the aggregate address must be in the BGP routing table for the configured aggregate to become active. If routes within the aggregate are constantly changing, do not add the `as-set` parameter to the aggregate because the aggregate flaps to track changes in the AS_PATH. The `no` version of this command disables the aggregate-address configuration.
Example	`OS10(conf-router-bgpv4-af)# aggregate-address 6.1.0.0/16 summary-only`
Supported Releases	10.3.0E or later

allowas-in

Configures the number of times the local AS number can appear in the BGP AS_PATH path attribute before the switch rejects the route.

Syntax	`allowas-in as-number`
Parameters	*as-number*—Enter the number of occurrences for a local AS number, from 1 to 10.
Default	Disabled
Command Mode	ROUTER-BPG-NEIGHBOR-AF
Usage Information	Use this command to enable the BGP speaker to accept a route with the local AS number in updates received from a peer for the specified number of times. The `no` version of this command resets the value to the default.
Example (IPv4)	`OS10(config-router-neighbor)# address-family ipv4 unicast` `OS10(conf-router-bgp-neighbor-af)# allowas-in 5`
Example (IPv6)	`OS10(conf-router-template)# address-family ipv6 unicast` `OS10(conf-router-bgp-template-af)# allowas-in 5`
Supported Releases	10.3.0E or later

always-compare-med

Compares MULTI_EXIT_DISC (MED) attributes in the paths received from different neighbors.

Syntax	`always-compare-med`
Parameters	None
Default	Disabled
Command Mode	ROUTER-BGP
Usage Information	After you use this command, use the `clear ip bgp *` command to recompute the best path. The `no` version of this command resets the value to the default. **NOTE: To configure these settings for a non-default VRF instance, first enter the ROUTER-CONFIG-VRF sub mode using the following commands:** 1 Enter the ROUTER BGP mode using the `router bgp as-number` command. 2 From the ROUTER BGP mode, enter ROUTER BGP VRF mode using the `vrf vrf-name` command.
Example	`OS10(conf-router-bgp-10)# always-compare-med`
Supported Releases	10.2.0E or later

as-notation

Changes the AS number notation format and requires four-octet-assupport.

Syntax	`as-format {asdot | asdot+ | asplain}`
Parameters	• `asdot` — Specify the AS number notation in asdot format. • `asdot+` — Specify the AS number notation in asdot+ format.

- `asplain` — Specify the AS number notation in asplain format.

Defaults `asplain`

Command Modes ROUTER-BGP

Usage Information

NOTE: To configure these settings for a non-default VRF instance, first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1 Enter the ROUTER BGP mode using the `router bgp` *`as-number`* command.
2 From the ROUTER BGP mode, enter ROUTER BGP VRF mode using the `vrf` *`vrf-name`* command.

Example

```
OS10(conf-router-bgp-2)# as-notation asdot
OS10(conf-router-bgp-2)# as-notation asdot+
OS10(conf-router-bgp-2)# as-notation asplain
```

Supported Releases 10.1.0E or later

bestpath as-path

Configures the AS path selection criteria for best path computation.

Syntax `bestpath as-path {ignore | mutlipath-relax}`

Parameters

- `ignore` — Enter to ignore the AS PATH in BGP best path calculations.
- `mutlipath-relax` — Enter to include prefixes received from different AS paths during multipath calculation.

Default Enabled

Command Mode ROUTER-BGP

Usage Information To enable load-balancing across different EBGP peers, configure the `mutlipath-relax` option. If you configure both `ignore` or `multipath-relax` options at the same time, a system-generated error message appears. The `no` version of this command disables configuration.

NOTE: To configure these settings for a non-default VRF instance, first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1 Enter the ROUTER BGP mode using the `router bgp` *`as-number`* command.
2 From the ROUTER BGP mode, enter ROUTER BGP VRF mode using the `vrf` *`vrf-name`* command.

Example

```
OS10(conf-router-bgp-10)# bestpath as-path multipath-relax
```

Supported Releases 10.3.0E or later

bestpath med

Changes the best path MED attributes during MED comparison for path selection.

Syntax `bestpath med {confed | missing-as-worst}`

Parameters

- `confed` — Compare MED among BGP confederation paths.

- `missing-as-worst` — Treat missing MED as the least preferred path.

Default Disabled

Command Mode ROUTER-BGP

Usage Information Before you apply this command, use the `always-compare-med` command. The `no` version of this command resets the MED comparison influence.

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp` *`as-number`* command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf` *`vrf-name`* command.

Example

```
OS10(conf-router-bgp-2)# bestpath med confed
```

Supported Releases 10.3.0E or later

bestpath router-id

Ignores comparing router-id information for external paths during best-path selection.

Syntax `bestpath router-id {ignore}`

Parameters `ignore` — Enter to ignore AS path for best-path computation.

Default Enabled

Command Mode ROUTER-BGP

Usage Information If you do not receive the same router ID for multiple paths, select the path that you received first. If you received the same router ID for multiple paths, ignore the path information. The `no` version of this command resets the value to the default.

NOTE: To configure these settings for a non-default VRF instance, first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp` *`as-number`* command.
2. From the ROUTER BGP mode, enter ROUTER BGP VRF mode using the `vrf` *`vrf-name`* command.

Example

```
OS10(conf-router-bgp-2)# bestpath router-id ignore
```

Supported Releases 10.3.0E or later

clear ip bgp

Resets BGP IPv4 or IPv6 neighbor sessions.

Syntax `clear ip bgp [vrf` *`vrf-name`*`] {`*`ipv4-address`* `|` *`ipv6-address`* `| * }`

Parameters

- `vrf` *`vrf-name`* — (OPTIONAL) Enter the keyword `vrf` then the name of the VRF to clear IPv4 or IPv6 BGP neighbor sessions corresponding to that VRF.
- *`IPv4-address`* — Enter an IPv4 address to clear a BGP neighbor configuration.
- *`IPv6-address`* — Enter an IPv6 address to clear a BGP neighbor configuration.

- `*` — Clears all BGP sessions.

Default	Not configured
Command Mode	EXEC
Usage Information	None.
Example	`OS10# clear ip bgp 1.1.15.4`
Supported Releases	10.3.0E or later

clear ip bgp *

Resets BGP sessions. The soft parameter, BGP soft reconfiguration, clears policies without resetting the TCP connection.

Syntax `clear ip bgp * [vrf vrf-name] [ipv4 unicast | ipv6 unicast | soft [in | out]]`

Parameters

- `*` — Enter to clear all BGP sessions.
- `vrf vrf-name` — (OPTIONAL) Enter the `vrf` then the name of the VRF to clear BGP session information corresponding to that VRF.
- `ipv4 unicast` — Enter to clear IPv4 unicast configuration.
- `ipv6 unicast` — Enter to clear IPv6 unicast configuration.
- `soft` — (Optional) Enter to configure and activate policies without resetting the BGP TCP session — BGP soft reconfiguration.
- `in` — (Optional) Enter to activate only ingress (inbound) policies.
- `out` — (Optional) Enter to activate only egress (outbound) policies.

Default	Not configured
Command Mode	EXEC
Usage Information	None.
Example	`OS10# clear ip bgp * ipv6 unicast`
Supported Releases	10.3.0E or later

clear ip bgp dampening

Clears the path information of the dampened and undampened prefixes.

Syntax `clear ip bgp dampening [vrf vrf-name] [ipv4-prefix | ipv6-prefix]`

Parameters

- `vrf vrf-name` — (OPTIONAL) Enter `vrf` then the name of the VRF to clear dampened paths information.
- `ipv4-prefix` — (Optional) Enter an IPv4 prefix of the dampened path.
- `ipv6-prefix` — (Optional) Enter an IPv6 prefix of the dampened path.

Default	Not configured
Command Mode	EXEC
Usage Information	None

Example	`OS10# clear ip bgp dampening 1.1.15.5`
Supported Releases	10.3.0E or later

clear ip bgp flap-statistics

Clears all or specific IPv4 or IPv6 flap counts of prefixes.

Syntax	`clear ip bgp [vrf vrf-name] [ipv4-address | ipv6-address] flap-statistics [ipv4-prefix | ipv6-prefix]`
Parameters	• `vrf vrf-name` — (OPTIONAL) Enter `vrf` then the name of the VRF to clear flap statistics information. • *ipv4-address* — (Optional) Enter an IPv4 address to clear the flap counts of the prefixes learned from the given peer. • *ipv6-address* — (Optional) Enter an IPv6 address to clear the flap counts. • *ipv4-prefix* — (Optional) Enter an IPv4 prefix to clear the flap counts of the given prefix. • *ipv6-prefix* — (Optional) Enter an IPv6 prefix to clear the flap counts of the given prefix.
Default	Not configured
Command Mode	EXEC
Usage Information	None
Example (All Prefixes)	`OS10# clear ip bgp flap-statistics`
Example (IPv4)	`OS10# clear ip bgp 1.1.15.4 flap-statistics`
Example (Given Prefix)	`OS10# clear ip bgp flap-statistics 1.1.15.0/24`
Supported Releases	10.3.0E or later

connection-retry-timer

Configures the timer to retry the connection to BGP neighbor or peer group.

Syntax	`connection-retry-timer retry-timer-value`
Parameters	*retry-timer-value* — Enter the time interval in seconds, ranging from 10 to 65535.
Defaults	60 seconds
Command Modes	CONFIG-ROUTER-NEIGHBOR CONFIG-ROUTER-TEMPLATE
Usage Information	The `no` version of this command resets the timer to default value..
Example	`OS10(config-router-neighbor)# connection-retry-timer 1000` `OS10(config-router-template)# connection-retry-timer 100`
Supported Releases	10.3.0E or later

confederation

Configures an identifier for a BGP confederation.

Syntax `confederation {identifier as-num | peers as-number}`

Parameters

- `identifier as-num` —Enter an AS number, from 0 to 65535 for 2 bytes, 1 to 4294967295 for 4 bytes, or 0.1 to 65535.65535 for dotted format.
- `peers as-number`—Enter an AS number for peers in the BGP confederation, from 1 to 4294967295.

Default Not configured

Command Mode ROUTER-BGP

Usage Information Configure your system to accept 4-byte formats before entering a 4-byte AS number. All routers in the Confederation must be 4-byte or 2-byte identified routers. You cannot have a mix of 2-byte and 4-byte identified routers. The autonomous system number you configure in this command is visible to the EBGP neighbors. Each autonomous system is fully meshed and contains a few connections to other autonomous systems. The next-hop (MED) and local preference information is preserved throughout the confederation. The system accepts confederation EBGP peers without a LOCAL_PREF attribute. OS10 sends AS_CONFED_SET and accepts AS_CONFED_SET and AS_CONF_SEQ. The `no` version of this command deletes the confederation configuration.

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp as-number` command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf vrf-name` command.

Example (Identifier)

```
OS10(conf-router-bgp-2)# confederation identifier 1
```

Example (Peers)

```
OS10(conf-router-bgp-2)# confederation peers 2
```

Supported Releases 10.3.0E or later

client-to-client

Enables route reflection between clients in a cluster.

Syntax `client-to-client {reflection}`

Parameters `reflection` — Enter to enable reflection of routes allowed in a cluster.

Default Enabled

Command Mode ROUTER-BGP

Usage Information Configure the route reflector to enable route reflection between all clients. You must fully mesh all clients before you disable route reflection. The `no` version of this command disables route reflection in a cluster.

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp as-number` command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf vrf-name` command.

Example

```
OS10(conf-router-bgp-2)# client-to-client reflection
```

Supported Releases 10.2.0E or later

cluster-id

Assigns a cluster ID to a BGP cluster with multiple route reflectors.

Syntax `cluster-id {number | ip-address}`

Parameters

- *number*—Enter a route reflector cluster ID as a 32-bit number, from 1 to 4294967295.
- *ip-address*—Enter an IP address as the route-reflector cluster ID.

Default Router ID

Command Mode ROUTER-BGP

Usage Information If a cluster contains only one route reflector, the cluster ID is the route reflector's router ID. For redundancy, a BGP cluster may contain two or more route reflectors. Without a cluster ID, the route reflector cannot recognize route updates from the other route reflectors within the cluster. The default format to display the cluster ID is A.B.C.D format. If you enter the cluster ID as an integer, an integer displays. The `no` version of this command resets the value to the default.

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp as-number` command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf vrf-name` command.

Example

```
OS10(conf-router-bgp-10)# cluster-id 3.3.3.3
```

Supported Releases 10.3.0E or later

bgp dampening

Enables BGP route-flap dampening and configures the dampening parameters.

Syntax `bgp dampening [half-life | reuse-limit | suppress-limit | max-suppress-time | route-map-name]`

Parameters

- *half-life* — (Optional) Enter the half-life time, in minutes, after which the penalty decreases. After the router assigns a penalty of 1024 to a route, the penalty decreases by half after the half-life period expires, from 1 to 45.
- *reuse-limit* — (Optional) Enter a reuse-limit value, which compares to the flapping route's penalty value. If the penalty value is less than the reuse value, the flapping route advertises again and is not suppressed, from 1 to 20000.

- *suppress-limit* — (Optional) Enter a suppress-limit value, which compares to the flapping route's penalty value. If the penalty value is greater than the suppress value, the flapping route is no longer advertised, from 1 to 20000.
- *max-suppress-time* — (Optional) Enter the maximum number of minutes a route is suppressed, from 1 to 255.
- *route-map-name* — (Optional) Enter the name of the route-map.

Defaults `half-life` 15; `reuse-limit` 750; `suppress-limit` 2000; `max-suppress-time` 60

Command Mode ROUTER-BGP-AF

Usage Information To reduce the instability of the BGP process, setup route flap dampening parameters. After setting up the dampening parameters, clear information on route dampening and return the suppressed routes to the `Active` state. You can also view statistics on route flapping or change the path selection from Default Deterministic mode to Non-Deterministic mode. The `no` version of this command resets the value to the default.

Example

```
OS10(conf-router-bgpv4-af)# dampening 2 751 2001 51 map1
```

Supported Releases 10.3.0E or later

description

Configures a description for the BGP neighbor or for peer template.

Syntax `description text`

Parameters *text* — Enter a description for the BGP neighbor or peer template.

Default None

Command Mode ROUTER-BGP-NEIGHBOR

ROUTER-BGP-TEMPLATE

Usage Information The `no` version of this command removes the description.

Example

```
OS10# configure terminal
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# neighbor 8.8.8.8
OS10(config-router-neighbor)# description  n1_abcd
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# template t1
OS10(config-router-template)# description peer_template_1_abcd
```

Supported Releases 10.4.1.0 or later

default-metric

Assigns a default-metric of redistributed routes to locally originated routes.

Syntax `default-metric number`

Parameters *number* — Enter a number as the metric to assign to routes from other protocols, from 1 to 4294967295.

Default Disabled

Command Mode ROUTER-BGP

Usage Information Assigns a metric for locally-originated routes such as redistributed routes. After you redistribute routes in BGP, use this command to reset the metric value — the new metric does not immediately take effect. The new metric takes effect only after you disable and re-enable route redistribution for a specified protocol. To re-enable route distribution use the `redistribute {connected [route-map map-name] | ospf process-id | static [route-map map-name]}` command, or use the `clear ip bgp *` command after you reset BGP. The `no` version of this command removes the default metric value.

Example (IPv4)

```
OS10(conf-router-bgpv4-af)# default-metric 60
```

Example (IPv6)

```
OS10(conf-router-bgpv6-af)# default-metric 60
```

Supported Releases 10.3.0E or later

default-originate

Configures the default route to a BGP peer or neighbor.

Syntax `default—originate [route-map route-map-name]`

Parameters `route-map route-map-name`—(Optional) Enter a route-map name. A maximum of 140 characters.

Default Enabled

Command Mode ROUTER-BGP-NEIGHBOR-AF

ROUTER-TEMPLATE-AF

Usage Information The `no` version of this command removes the default route.

Example

```
OS10(conf-router-bgp-10)# template lunar
OS10(conf-router-bgp-template)# address-family ipv6 unicast
OS10(conf-router-template-af)# default-originate route-map rmap-bgp
```

Supported Releases 10.4.1.0 or later

distance bgp

Sets the administrative distance for BGP routes.

Syntax `distance bgp external-distance internal-distance local-distance`

Parameters

- *external-distance*—Enter a number to assign to routes learned from a neighbor external to the AS, from 1 to 255.
- *internal-distance*—Enter a number to assign to routes learned from a router within the AS, from 1 to 255.
- *local-distance*—Enter a number to assign to routes learned from networks listed in the `network` command, from 1 to 255.

Defaults

- *external-distance*—20
- *internal-distance*—200
- *local-distance*—200

Command Modes

- CONFIG-ROUTER-BGP-ADDRESS-FAMILY
- CONFIG-ROUTER-BGP-VRF-ADDRESS-FAMILY

Usage Information

This command is used to configure administrative distance for eBGP route, iBGP route, and local BGP route. Administrative distance indicates the reliability of the route; the lower the administrative distance, the more reliable the route is. Routes that are assigned an administrative distance of 255 are not installed in the routing table. Routes from confederations are treated as iBGP routes.

Examples

Default VRF:

IPv4

```
OS10# configure terminal
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# address-family ipv4 unicast
OS10(configure-router-bgpv4-af)# distance bgp 10 200 210
```

IPv6

```
OS10# configure terminal
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# address-family ipv6 unicast
OS10(configure-router-bgpv6-af)# distance bgp 10 200 210
```

Non-default VRF

```
OS10(config-router-bgp-100)# vrf blue
OS10(config-router-bgp-100-vrf)# address-family ipv4 unicast
OS10(configure-router-bgpv4-vrf-af)# distance bgp 21 200 200
OS10(config-router-bgp-100-vrf)# address-family ipv6 unicast
OS10(configure-router-bgpv6-vrf-af)# distance bgp 21 201 250
```

Supported Releases

10.4.2.0 or later

distribute-list

Distributes BGP information through an established prefix list.

Syntax

`distribute-list` *`prefix-list-name`* `{in | out}`

Parameters

- *`prefix-list-name`*—Enter the name of established prefix list.
- `in`—Enter to distribute inbound traffic.
- `out`—Enter to distribute outbound traffic.

Defaults

None

Command Modes

ROUTER-BGP-NEIGHBOR-AF

ROUTER-TEMPLATE-AF

Usage Information

The `no` version of this command removes the `route-map`.

Example

```
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-bgp-neighbor-af)# distribute-list inbgg in
```

```
OS10(conf-router-template)# address-family ipv4 unicast
OS10(conf-router-bgp-template-af)# distribute-list outbgg out
```

Supported Releases	10.4.1.0 or later

bgp default local-preference

Changes the default local preference value for routes exchanged between internal BGP peers.

Syntax	`default local-preference number`
Parameters	*number* — Enter a number to assign to routes as the degree of preference for those routes. When routes compare, the route with the higher degree of preference or the local preference value is most preferred, from 1 to 4294967295.
Default	100
Command Mode	ROUTER-BGP
Usage Information	All routers apply this command setting within the AS. The `no` version of this command removes local preference value.
Example	`OS10(conf-router-bgp-1)# default local-preference 200`
Supported Releases	10.3.0E or later

ebgp-multihop

Allows EBGP neighbors on indirectly connected networks.

Syntax	`ebgp-multihop hop count`
Parameters	*hop count* — Enter a value for the number of hops, from 1 to 255.
Default	1
Command Mode	ROUTER-NEIGHBOR
Usage Information	This command avoids installation of default multihop peer routes to prevent loops and creates neighbor relationships between peers. Networks indirectly connected are not valid for best path selection. The `no` version of this command removes multihop session.
Example	`OS10(conf-router-neighbor)# ebgp-multihop 2`
Supported Releases	10.3.0E or later

enforce-first-as

Enforces the first AS in the AS path of the route received from an EBGP peer to be the same as the configured remote AS.

Syntax	`enforce-first-as`
Parameters	None
Default	Enabled
Command Mode	ROUTER-BGP
Usage Information	To verify statistics of routes rejected, use the `show ip bgp neighbors` command. If routes are rejected, the session is reset. In the event of a failure, the existing BGP sessions flap. For updates received from EBGP peers,

BGP ensures that the first AS of the first AS segment is always the AS of the peer, otherwise the update drops and the counter increments. The `no` version of this command turns off the default.

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp` *`as-number`* command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf` *`vrf-name`* command.

Example

```
OS10(conf-router-bgp-1)# enforce-first-as
```

Supported Releases 10.3.0E or later

fall-over

Enables or disables BGP session fast fall-over for BGP neighbors.

Syntax `fall-over`

Parameters None

Default Disabled

Command Mode ROUTER-NEIGHBOR

Usage Information Configure the BGP fast fall-over on a per-neighbor or peer-group basis. When you enable this command on a template, it simultaneously enables on all peers that inherit the peer group template. When you enable fall-over, BGP tracks IP reachability to the peer remote address and the peer local address. Whenever either address becomes unreachable — no active route exists in the routing table for peer IPv6 destinations or local address — BGP brings down the session with the peer. The `no` version of this command disables fall-over.

Example

```
OS10(conf-router-neighbor)# fall-over
```

Supported Releases 10.3.0E or later

fast-external-fallover

Resets BGP sessions immediately when a link to a directly connected external peer fails.

Syntax `fast-external-fallover`

Parameters None

Default Not configured

Command Mode ROUTER-BGP

Usage Information Fast external fall-over terminates the EBGP session immediately after the IP unreachability or link failure is detected. This only applies after you manually reset all existing BGP sessions. For the configuration to take effect, use the `clear ip bgp` command. The `no` version of this command disables fast external fallover.

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp` *`as-number`* command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf` *`vrf-name`* command.

Example	`OS10(conf-router-bgp-10)# fast-external-fallover`
Supported Releases	10.3.0E or later

inherit template

Configures a peer group template name that the neighbors use to inherit peer-group configuration.

Syntax	`inherit template` *`template-name`*
Parameters	*`template-name`* — Enter a template name. A maximum of 16 characters.
Default	Not configured
Command Mode	ROUTER-NEIGHBOR
Usage Information	When network neighbors inherit a template, all features enabled on the template are also supported on the neighbors. The `no` version of this command disables the peer group template configuration.
Example	`OS10(conf-router-neighbor)# inherit template zanzibar`
Supported Releases	10.2.0E or later

listen

Enables peer listening and sets the prefix range for dynamic peers.

Syntax	`listen` *`ip-address`* `[limit` *`count`*`]`
Parameters	• *`ip-address`*—Enter the BGP neighbor IP address. • `limit` *`count`*—(Optional) Enter a maximum dynamic peer count, from 1 to 4294967295.
Default	Not configured
Command Mode	ROUTER-TEMPLATE
Usage Information	Enables a passive peering session for listening. The `no` version of this command disables a passive peering session.
Example	`OS10(conf-router-template)# listen 1.1.0.0/16 limit 4`
Supported Releases	10.2.0E or later

local-as

Configures a local AS number for a peer.

Syntax	`local-as` *`as-number`* `[no-prepend]`
Parameters	• *`as-number`*—Enter the local AS number, from 1 to 4294967295. • `no-prepend`—(Optional) Enter so that local AS values are not prepended to announcements from the neighbor.
Default	Disabled
Command Mode	ROUTER-NEIGHBOR or ROUTER-TEMPLATE

Usage Information	Facilitates the BGP network migration operation and allows you to maintain existing AS numbers. The `no` version of this command resets the value to the default.
Example (Neighbor)	`OS10(conf-router-bgp-10)# neighbor lunar` `OS10(conf-router-neighbor)# local-as 20`
Example (Template)	`OS10(conf-router-bgp-10)# template solar` `OS10(conf-router-template)# local-as 20`
Supported Releases	10.3.0E or later

log-neighbor-changes

Enables logging for changes in neighbor status.

Syntax	`log-neighbor-changes`
Parameters	None
Default	Enabled
Command Mode	ROUTER-BGP
Usage Information	OS10 saves logs which includes the neighbor operational status and reset reasons. To view the logs, use the `show bgp config` command. The `no` version of this command disables the feature.
Example	`OS10(conf-router-bgp-10)# log-neighbor-changes`
Supported Releases	10.3.0E or later

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp` *`as-number`* command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf` *`vrf-name`* command.

maximum-paths

Configures the maximum number of equal-cost paths for load sharing.

Syntax	`maximum-paths [ebgp` *`number`* `| ibgp` *`number`*`] maxpaths`
Parameters	• `ebgp`—Enable multipath support for external BGP routes. • `ibgp`—Enable multipath support for internal BGP routes. • *`number`*—Enter the number of parallel paths, from 1 to 64.
Default	64 paths
Command Mode	ROUTER-BGP
Usage Information	Dell EMC recommends not using multipath and add path simultaneously in a route reflector. To recompute the best path, use the `clear ip bgp *` command. The `no` version of this command resets the value to the default

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp as-number` command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf vrf-name` command.

Example (EBGP)
```
OS10(conf-router-bgp-2)# maximum-paths ebgp 2 maxpaths
```

Example (IBGP)
```
OS10(conf-router-bgp-2)# maximum-paths ibgp 4 maxpaths
```

Supported Releases 10.3.0E or later

maximum-prefix

Configures the maximum number of prefixes allowed from a peer.

Syntax
```
maximum-prefix {number [threshold] [warning]}
```

Parameters

- *number*—Enter a maximum prefix number, from 1 to 4294967295.
- *threshold*—(Optional) Enter a threshold percentage, from 1 to 100.
- `warning-only` — (Optional) Enter to set the router to send a warning log message when the maximum limit is exceeded. If you do not set this parameter, the router stops peering when the maximum prefixes limit exceeds.

Default 75% threshold

Command Mode ROUTER-BGP-NEIGHBOR-AF

Usage Information If you configure this command and the neighbor receives more prefixes than the configuration allows, the neighbor goes down. To view the prefix information, use the `show ip bgp summary` command. The neighbor remains down until you use the `clear ip bgp` command for the neighbor or the peer group to which the neighbor belongs. The `no` version of this command resets the value to the default.

Example
```
OS10(conf-router-bgp-neighbor-af)# maximum-prefix 20 100 warning-only
```

Supported Releases 10.3.0E or later

neighbor

Creates a remote peer for the BGP neighbor and enters BGP Neighbor mode.

Syntax
```
neighbor ip address
```

Parameters *ip address* — Enter the IP address of the neighbor in dotted decimal format.

Default Not configured

Command Mode CONFIG-ROUTER-BGP

Usage Information Create a remote peer with the BGP neighbor. Always enter the IP address of a BGP peer with this command. The command does not validate if the configured peer address is a local IP address. The `no` version of this command disables the BGP neighbor configuration.

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp as-number` command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf vrf-name` command.

Example

```
OS10(conf-router-bgp-2)# neighbor 32.1.0.0
OS10(conf-router-neighbor)#
```

Supported Releases 10.3.0E or later

next-hop-self

Disables the next-hop calculation for a neighbor.

Syntax `next-hop-self`

Parameters None

Default Enabled

Command Mode ROUTER-NEIGHBOR-AF

Usage Information Influences next-hop processing of EBGP routes to IBGP peers. The `no` version of this command disables the next-hop calculation.

Example

```
OS10(conf-router-neighbor-af)# next-hop-self
```

Supported Releases 10.3.0E or later

non-deterministic-med

Compares paths in the order they arrive.

Syntax `non-deterministic-med`

Parameters None

Default Disabled

Command Mode ROUTER-BGP

Usage Information Paths compare in the order they arrive. OS10 uses this method to choose different best paths from a set of paths, depending on the order they are received from the neighbors. MED may or may not be compared between adjacent paths. When you change the path selection from deterministic to non-deterministic, the path selection for the existing paths remains deterministic until you use the `clear ip bgp` command to clear the existing paths. The `no` version of this command configures BGP bestpath selection as non-deterministic.

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp as-number` command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf vrf-name` command.

Example

```
OS10(conf-router-bgp-10)# non-deterministic-med
```

Supported Releases 10.2.0E or later

outbound-optimization

Enables outbound optimization for IBGP peer-group members.

Syntax `outbound-optimization`

Parameters None

Default Not configured

Command Mode ROUTER-BGP

Usage Information Enable or disable outbound optimization dynamically to reset all neighbor sessions. When you enable outbound optimization, all peers receive the same update packets. The next-hop address chosen as one of the addresses of neighbor's reachable interfaces is also the same for the peers. The `no` version of this command disables outbound optimization.

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp` *`as-number`* command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf` *`vrf-name`* command.

Example

```
OS10(conf-router-bgp-10)# outbound-optimization
```

Supported Releases 10.3.0E or later

password

Configures a password for message digest 5 (MD5) authentication on the TCP connection between two neighbors.

Syntax `password {9` *`encrypted password-string`*`| password-string}`

Parameters

- `9` *`encrypted password-string`*—Enter 9 then the encrypted password.
- *`password-string`*—Enter a password for authentication. A maximum of 128 characters.

Default Disabled

Command Mode ROUTER-NEIGHBOR

ROUTER-TEMPLATE

Usage Information You can enter the password either as plain text or in encrypted format. The password provided in ROUTER-NEIGHBOR mode takes preference over the password in ROUTER-TEMPLATE mode. The `no` version of this command disables authentication.

Example

```
OS10(conf-router-neighbor)# password abcdell
```

```
OS10(conf-router-neighbor)# password 9
f785498c228f365898c0efdc2f476b4b27c47d972c3cd8cd9b91f518c14ee42d
```

Supported Releases 10.3.0E or later

redistribute

Redistributes connected, static, and OSPF routes in BGP.

Syntax

```
redistribute {connected [route-map map name] | ospf process-id | static [route-map map name ]}
```

Parameters

- `connected` — Enter to redistribute routes from physically connected interfaces.
- `route-map map name` — (Optional) Enter the name of a configured route-map.
- `ospf process-id` — Enter a number for the OSPF process (1 to 65535).
- `static` — Enter to redistribute manually configured routes.

Default Disabled

Command Mode ROUTER-BGPv4-AF or ROUTER-BGPv6-AF

Usage Information Static routes are treated as incomplete routes. When you use the `redistribute ospf process-id` command without other parameters, the system redistributes all OSPF internal routes, external type 1 routes, and external type 2 routes. The `no` version of this command resets the value to the default.

Example (Connected)

```
OS10(conf-router-bgp-102)# address-family ipv4 unicast
OS10(conf-router-bgpv4-af)# redistribute connected route-map mapbgp1
```

Example (Static — IPv4)

```
OS10(conf-router-bgp-102)# address-family ipv4 unicast
OS10(conf-router-bgpv4-af)# redistribute static route-map mapbgp2
```

Example (Static — IPv6)

```
OS10(conf-router-bgp-102)# address-family ipv6 unicast
OS10(conf-router-bgpv6-af)# redistribute static
```

Example (OSPF — IPv4)

```
OS10(conf-router-bgp-102)# address-family ipv4 unicast
OS10(conf-router-bgpv4-af)# redistribute ospf 1
```

Example (OSPF — IPv6)

```
OS10(conf-router-bgp-102)# address-family ipv6 unicast
OS10(conf-router-bgpv6-af)# redistribute ospf 1
```

Supported Releases 10.2.0E or later

remote-as

Adds a remote AS to the specified BGP neighbor or peer group.

Syntax

```
remote-as as-number
```

Parameters *as-number* — Specify AS number ranging from 1 to 65535 for 2-byte or 1 to 4294967295 for 4-byte.

Defaults None

Command Modes CONFIG-ROUTER-NEIGHBOR

CONFIG-ROUTER-TEMPLATE

Usage Information The `no` version of this command removes the remote AS.

Example

```
OS10(config)# router bgp 300
OS10(config-router-bgp-300)# template ebgppg
OS10(config-router-template)# remote-as 100
```

Supported Releases 10.4.1.0 or later

remove-private-as

Removes private AS numbers from receiving outgoing updates.

Syntax `remove-private-as`

Parameters None

Defaults Disabled

Command Mode CONFIG-ROUTER-NEIGHBOR

CONFIG-ROUTER-TEMPLATE

Usage Information None

Example

```
OS10(config)# router bgp 300
OS10(config-router-bgp-300)# template ebgppg
OS10(config-router-template)# remove-private-as
```

Supported Releases 10.4.1.0 or later

route-map

Applies an established route-map to either incoming or outbound routes of a BGP neighbor or peer group.

Syntax `route-map` *`route-map-name`* `{in | out}`

Parameters

- *`route-map-name`* — Enter the name of the configured route-map.
- `in` — attaches the route-map as the inbound policy
- `out`— attaches the route-map as the outbound policy

Defaults None

Command Modes ROUTER-BGP-TEMPLATE-AF

Usage Information The `no` version of this command removes the `route-map`.

Example

```
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-bgp-neighbor-af)# route-map bgproutemap in
```

```
OS10(conf-router-template)# address-family ipv4 unicast
OS10(conf-router-bgp-template-af)# route-map bgproutemap in
```

Supported Releases 10.4.1.0 or later

route-reflector-client

Configures a neighbor as a member of a route-reflector cluster.

Syntax	`route-reflector-client`
Parameters	None
Default	Not configured
Command Mode	ROUTER-TEMPLATE
Usage Information	The device configures as a route reflector, and the BGP neighbors configure as clients in the route-reflector cluster. The `no` version of this command removes all clients of a route reflector—the router no longer functions as a route reflector.
Example	`OS10(conf-router-template)# route-reflector-client`
Supported Releases	10.3.0E or later

router bgp

Enables BGP and assigns an AS number to the local BGP speaker.

Syntax	`router bgp` *`as-number`*
Parameters	*`as-number`*—Enter the AS number range. • 1 to 65535 in 2-byte • 1 to 4294967295 in 4-byte
Default	None
Command Mode	CONFIGURATION
Usage Information	The AS number can be a 16-bit integer. The `no` version of this command resets the value to the default.
Example	`OS10(config)# router bgp 3` `OS10(conf-router-bgp-3)#`
Supported Releases	10.3.0E or later

router-id

Assigns a user-given ID to a BGP router.

Syntax	`router-id` *`ip-address`*
Parameters	*`ip-address`* — Enter an IP address in dotted decimal format.
Default	First configured IP address or random number
Command Mode	ROUTER-BGP
Usage Information	Change the router ID of a BGP router to reset peer-sessions. The `no` version of this command resets the value to the default.

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp as-number` command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf vrf-name` command.

Example

```
OS10(conf-router-bgp-10)# router-id 10.10.10.40
```

Supported Releases 10.3.0E or later

send-community

Sends a community attribute to a BGP neighbor or peer group.

Syntax `send-community {extended | standard}`

Parameters

- `extended` — Enter an extended community attribute.
- `standard` — Enter a started community attribute.

Default Not configured

Command Mode ROUTER-NEIGHBOR

Usage Information A community attribute indicates that all routes with the same attributes belong to the same community grouping. All neighbors belonging to the template inherit the feature when configured for a template. The `no` version of this command disables sending a community attribute to a BGP neighbor or peer group.

Example

```
OS10(conf-router-neighbor)# send-community extended
```

Supported Releases 10.3.0E or later

sender-side-loop-detection

Enables the sender-side loop detection process for a BGP neighbor.

Syntax `sender-side-loop-detection`

Parameters None

Default Enabled

Command Mode ROUTER-BGP-NEIGHBOR-AF

Usage Information This command helps detect routing loops, based on the AS path before it starts advertising routes. To configure a neighbor to accept routes use the `neighbor allowas-in` command. The `no` version of this command disables sender-side loop detection for that neighbor.

Example (IPv4)

```
OS10(conf-router-bgp-102)# neighbor 3.3.3.1
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-bgp-neighbor-af)# sender-side-loop-detection
```

Example (IPv6)

```
OS10(conf-router-bgp-102)# neighbor 32::1
OS10(conf-router-neighbor)# address-family ipv6 unicast
OS10(conf-router-bgp-neighbor-af)# no sender-side-loop-detection
```

Supported Releases 10.3.0E or later

show ip bgp

Displays information that BGP neighbors exchange.

Syntax

```
show ip bgp [vrf vrf-name] ip-address/mask
```

Parameters

- `vrf vrf-name` — (OPTIONAL) Enter `vrf` and then the name of the VRF to view route information corresponding to that VRF.
- *`ip-address/mask`* — Enter the IP address and mask in A.B.C.D/x format.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip bgp 1.1.1.0/24
BGP routing table entry for 1.1.1.0/24
Paths: (1 available, table Default-IP-Routing-Table.)

Received from :
3.1.1.1(3.3.3.33) Best

AS_PATH : 100
Next-Hop : 3.1.1.1, Cost : 0

Origin INCOMPLETE, Metric 0, LocalPref 100, Weight  0, confed-external
Route-reflector origin : 0.0.0.0
```

Supported Releases 10.3.0E or later

show ip bgp dampened-paths

Displays BGP routes that are dampened or non-active.

Syntax

```
show ip bgp [vrf vrf-name] dampened-paths
```

Parameters None

Default Not configured

Command Mode EXEC

Usage Information

- `vrf vrf-name` — (OPTIONAL) Enter `vrf` and then the name of the VRF to view routes that are affected by a specific community list corresponding to that VRF.
- `Network` — Displays the network ID where the route is dampened.
- `From` — Displays the IP address of the neighbor advertising the dampened route.
- `Reuse` — Displays the HH:MM:SS until the dampened route is available.
- `Path` — Lists all AS the dampened route passed through to reach the destination network.

Example

```
OS10# show ip bgp dampened-paths

BGP local router ID is 80.1.1.1
Status codes: s suppressed, S stale, d dampened, h history, * valid, > best
Origin codes: i - IGP, e - EGP, ? - incomplete
          Network        From           Reuse          Path
d*        3.1.2.0/24     80.1.1.2       00:00:12       800 9 8 i
```

```
d*        3.1.3.0/24     80.1.1.2       00:00:12       800 9 8 i
d*        3.1.4.0/24     80.1.1.2       00:00:12       800 9 8 i
d*        3.1.5.0/24     80.1.1.2       00:00:12       800 9 8 i
d*        3.1.6.0/24     80.1.1.2       00:00:12       800 9 8 i
Total number of prefixes: 5
```

Supported Releases 10.3.0E or later

show ip bgp flap-statistics

Displays BGP flap statistics on BGP routes.

Syntax `show ip bgp [vrf vrf-name] flap-statistics`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information

- `vrf vrf-name` — (OPTIONAL) Enter `vrf` and then the name of the VRF to view flap statistics on BGP routes corresponding to that VRF.
- `Network` — Displays the network ID where the route is flapping.
- `From` — Displays the IP address of the neighbor advertising the flapping route.
- `Duration` — Displays the HH:MM:SS after the route first flapped.
- `Flaps` — Displays the number of times the route flapped.
- `Reuse` — Displays the HH:MM:SS until the flapped route is available.
- `Path` — Lists all AS the flapping route passed through to reach the destination network.

Example

```
OS10# show ip bgp flap-statistics
BGP local router ID is 80.1.1.1
Status codes: s suppressed, S stale, d dampened, h history, * valid, > best
Origin codes: i - IGP, e - EGP, ? - incomplete
     Network      From       Flaps  Duration  Reuse      Path
*>   3.1.2.0/24   80.1.1.2   1      00:00:11  00:00:00   800 9 8 i
*>   3.1.3.0/24   80.1.1.2   1      00:00:11  00:00:00   800 9 8 i
*>   3.1.4.0/24   80.1.1.2   1      00:00:11  00:00:00   800 9 8 i
*>   3.1.5.0/24   80.1.1.2   1      00:00:11  00:00:00   800 9 8 i
*>   3.1.6.0/24   80.1.1.2   1      00:00:11  00:00:00   800 9 8 i
Total number of prefixes: 5
```

Supported Releases 10.3.0E or later

show ip bgp ipv4 unicast

Displays route information for BGP IPv4 routes.

Syntax `show ip bgp [vrf vrf-name] ipv4 unicast {ip-address/mask | summary} [denied-routes]`

Parameters

- `vrf vrf-name` — (OPTIONAL) Enter `vrf` then the name of the VRF to view IPv4 unicast summary information corresponding to that VRF.
- `unicast ip-address/mask` — Displays IPv4 unicast route information.
- `summary` — Displays IPv4 unicast summary information.

- `denied-routes` — (Optional) Displays the configured denied routes.

Default Not configured

Command Mode EXEC

Usage Information This command displays locally advertised BGPv4 routes configured using the `network` command. These routes show as `r` for redistributed/network-learned routes.

Example

```
OS10# show ip bgp ipv4 unicast summary
BGP router identifier 80.1.1.1 local AS number 102
Neighbor     AS    MsgRcvd  MsgSent Up/Down   State/Pfx
80.1.1.2     800   8        4       00:01:10  5
```

Supported Releases 10.3.0E or later

show ip bgp ipv6 unicast

Displays route information for BGP IPv6 routes.

Syntax

```
show ip bgp [vrf vrf-name] ipv6 unicast [neighbors] {ip-address/mask | summary}
| multicast {ip-address/mask | neighbors} [denied-routes]
```

Parameters

- `vrf vrf-name` — (OPTIONAL) Enter `vrf` then the name of the VRF to view IPv6 unicast information corresponding to that VRF.
- `neighbors` — Displays IPv6 neighbor information.
- *`ip-address/mask`* — Displays information about IPv6 unicast routes.
- `summary` — Displays IPv6 unicast summary information.
- `multicast` *`ip-address/mask`* — Displays IPv6 multicast routes information.
- `denied-routes` — (Optional) Displays the configured IPv6 denied routes.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip bgp ipv6 unicast summary
BGP router identifier 80.1.1.1 local AS number 102
Neighbor    AS   MsgRcvd  MsgSent Up/Down    State/Pfx
80.1.1.2    800  8        4       00:01:10   5
```

Supported Releases 10.3.0E or later

show ip bgp neighbors

Displays information that BGP neighbors exchange.

Syntax

```
show ip bgp [vrf vrf-name] neighbors ip-address [denied-routes]
```

Parameters

- `vrf vrf-name` — (OPTIONAL) Enter `vrf` and then the name of the VRF to view information exchanged between BGP neighbors corresponding to that VRF
- *`ip-address`* — Enter the IP address for a specific neighbor
- `denied-routes` — (Optional) Displays the list of routes denied by policy
- `advertised-routes`—Displays the routes advertised to a neighbor

- `dampened-routes`—Displays the suppressed routes received from a neighbor
- `flap-statistics`—Displays the route's flap statistics received from a neighbor
- `received-routes`—Displays the routes received from a neighbor
- `routes`—Displays routes learned from a neighbor

Default Not configured

Command Mode EXEC

Usage Information

- `BGP neighbor` — Displays the BGP neighbor address and its AS number. The last phrase in the line indicates whether the link between the BGP router and its neighbor is an external or internal one. If they are located in the same AS, the link is internal; otherwise the link is external.
- `BGP version` — Displays the BGP version, always version 4, and the remote router ID.
- `BGP state` — Displays the neighbor's BGP state and the amount of time in hours:minutes: seconds it has been in that state.
- `Last read` — Displays the information included in the last read:
  - Last read is the time in hours:minutes:seconds that the router read a message from its neighbor.
  - Hold time is the number of seconds configured between messages from its neighbor.
  - Keepalive interval is the number of seconds between keepalive messages to help ensure that the TCP session is still alive.
- `Received messages` — Displays the number of BGP messages received, the number of notifications or error messages, and the number of messages waiting in a queue for processing.
- `Sent messages` — Displays the number of BGP messages sent, the number of notifications or error messages, and the number of messages waiting in a queue for processing.
- `Description` — Displays the descriptive name configured for the BGP neighbor. This field is displayed only when the description is configured.
- `Local host` — Displays the peering address of the local router and the TCP port number.
- `Foreign host` — Displays the peering address of the neighbor and the TCP port number.

Although the status codes for routes received from a BGP neighbor may not display in the `show ip bgp neighbors` *`ip-address`* `received-routes` output, they display correctly in the `show ip bgp` output.

Example

```
OS10# show ip bgp neighbors
BGP neighbor is 80.1.1.2, remote AS 800, local AS 102  external link

  BGP version 4, remote router ID 12.12.0.2
  BGP state ESTABLISHED, in this state for 00:02:51
  Last read 00:18:23 seconds
  Hold time is 90, keepalive interval is 30 seconds
  Configured hold time is 180, keepalive interval is 60 seconds
  Fall-over disabled

  Received 11 messages
    1 opens, 0 notifications, 3 updates
    7 keepalives, 0 route refresh requests
  Sent 8 messages
    1 opens, 0 notifications, 0 updates
    7 keepalives, 0 route refresh requests
  Minimum time between advertisement runs is 30 seconds
  Description: n1_abcd
  Capabilities received from neighbor for IPv4 Unicast:
MULTIPROTO_EXT(1)MULTIPROTO_EXT(1)MULTIPROTO_EXT(1)ROUTE_REFRESH(2)
  Capabilities advertised to neighbor for IPv4 Unicast:
MULTIPROTO_EXT(1)MULTIPROTO_EXT(1)ROUTE_REFRESH(2)CISCO_ROUTE_REFRESH
  (128)4_OCTET_AS(65)
  Prefixes accepted 5, Prefixes advertised 0
  Connections established 1; dropped 1
Closed by neighbor sent 00:02:51 ago
```

```
  For address family: IPv4 Unicast
  Next hop set to self
  Allow local AS number 0 times in AS-PATH attribute

  For address family: IPv6 Unicast
  Next hop set to self
  Allow local AS number 0 times in AS-PATH attribute

  Local host: 80.1.1.1, Local port: 57812
  Foreign host: 80.1.1.2, Foreign port: 179
```

Example advertised-routes

```
OS10# show ip bgp ipv6 unicast neighbors 192:168:1::2 advertised-routes
BGP local router ID is 100.1.1.1
Status codes: s suppressed, S stale, d dampened, h history, * valid, > best
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop       Metric    LocPrf    Weight  Path
*>55::/64        192:168:1::1   0         0         0       100i
*>55:0:0:1::/64  192:168:1::1   0         0         0       100i
*>55:0:0:2::/64  192:168:1::1   0         0         0       100i
*>55:0:0:3::/64  192:168:1::1   0         0         0       100i
*>55:0:0:4::/64  192:168:1::1   0         0         0       100i
*>55:0:0:5::/64  192:168:1::1   0         0         0       100i
*>55:0:0:6::/64  192:168:1::1   0         0         0       100i
*>55:0:0:7::/64  192:168:1::1   0         0         0       100i
*>55:0:0:8::/64  192:168:1::1   0         0         0       100i
*>55:0:0:9::/64  192:168:1::1   0         0         0       100i
*>172:16:1::/64  192:168:1::1   0         0         0       100?
Total number of prefixes: 11
OS10#
```

Example received-routes

```
OS10# show ip bgp ipv6 unicast neighbors 172:16:1::2 received-routes
BGP local router ID is 100.1.1.1
Status codes: D denied
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop       Metric    LocPrf    Path
D  55::/64        172:16:1::2    0         0         i
   55:0:0:1::/64  172:16:1::2    0         0         i
   55:0:0:2::/64  172:16:1::2    0         0         i
D  55:0:0:3::/64  172:16:1::2    0         0         i
D  55:0:0:4::/64  172:16:1::2    0         0         i
D  55:0:0:5::/64  172:16:1::2    0         0         i
D  55:0:0:6::/64  172:16:1::2    0         0         i
   55:0:0:7::/64  172:16:1::2    0         0         i
D  55:0:0:8::/64  172:16:1::2    0         0         i
D  55:0:0:9::/64  172:16:1::2    0         0         i
Total number of prefixes: 10
OS10#
```

Example denied-routes

```
OS10# show ip bgp ipv6 unicast neighbors 172:16:1::2 denied-routes
BGP local router ID is 100.1.1.1
Status codes: D denied
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop     Metric   LocPrf  Path
D  55::/64        172:16:1::2  0        0       100 200 300 400i
D  55:0:0:1::/64  172:16:1::2  0        0       100 200 300 400i
D  55:0:0:2::/64  172:16:1::2  0        0       100 200 300 400i
Total number of prefixes: 3
OS10#
```

Example routes

```
OS10# show ip bgp ipv6 unicast neighbors 172:16:1::2 routes
BGP local router ID is 100.1.1.1
Status codes: s suppressed, S stale, d dampened, h history, * valid, > best
Origin codes: i - IGP, e - EGP, ? - incomplete
Network             Next Hop       Metric    LocPrf   Weight  Path
*>55::/64           172:16:1::2    44        55       0       i
*>55:0:0:1::/64     172:16:1::2    44        55       0       i
```

```
*>55:0:0:2::/64     172:16:1::2    44         55        0        i
*>55:0:0:3::/64     172:16:1::2    44         55        0        i
*>55:0:0:4::/64     172:16:1::2    44         55        0        i
*>55:0:0:5::/64     172:16:1::2    44         55        0        i
*>55:0:0:6::/64     172:16:1::2    44         55        0        i
*>55:0:0:7::/64     172:16:1::2    44         55        0        i
*>55:0:0:8::/64     172:16:1::2    44         55        0        i
*>55:0:0:9::/64     172:16:1::2    44         55        0        i
Total number of prefixes: 10
OS10#
```

Supported Releases 10.3.0E or later

show ip bgp peer-group

Displays information on BGP peers in a peer-group.

Syntax `show ip bgp [vrf vrf-name] peer-group peer-group-name`

Parameters

- `vrf vrf-name` — (OPTIONAL) Enter `vrf` to view information on BGP peers in a peer group corresponding to that VRF.
- *`peer-group-name`* — (Optional) Enter the peer group name to view information about that peer-group only.

Default Not configured

Command Mode EXEC

Usage Information

- `Peer-group` — Displays the peer group name. Minimum time displays the time interval between BGP advertisements.
- `Administratively shut` — Displays the peer group's status if you do not enable the peer group. If you enable the peer group, this line does not display.
- `BGP version` — Displays the BGP version supported.
- `Description` — Displays the descriptive name configured for the BGP peer template. This field is displayed only when the description is configured.
- `For address family` — Displays IPv4 unicast as the address family.
- `BGP neighbor` — Displays the name of the BGP neighbor.
- Number of peers — Displays the number of peers currently configured for this peer group.
- `Peer-group members` — Lists the IP addresses of the peers in the peer group. If the address is outbound optimized, an * displays next to the IP address.

Example

```
OS10# show ip bgp peer-group bgppg
Peer-group bgppg, remote AS 103
   BGP version 4
   Minimum time between advertisement runs is 30 seconds
   Description: peer_template_1_abcd
   For address family: Unicast
   BGP neighbor is bgppg, peer-group external
   Update packing has 4_OCTET_AS support enabled
```

Example (Summary)

```
OS10# show ip bgp peer-group ebgp summary
BGP router identifier 32.1.1.1  local AS number 6
Neighbor      AS   MsgRcvd   MsgSent   Up/Down   State/Pfx
17.1.1.2      7    7         6         00:01:54  5
```

Supported Releases 10.2.0E or later

show ip bgp summary

Displays the status of all BGP connections.

Syntax `show ip bgp [vrf vrf-name] summary`

Parameters `vrf vrf-name` — (OPTIONAL) Enter `vrf` then the name of the VRF to view the status of all BGP connections corresponding to that VRF.

Default Not configured

Command Mode EXEC

Usage Information

- `Neighbor`—Displays the BGP neighbor address.
- `AS`—Displays the AS number of the neighbor
- `MsgRcvd`—Displays the number of BGP messages that the neighbor received.
- `MsgSent`—Displays the number of BGP messages that the neighbor sent.
- `Up/Down`—Displays the amount of time that the neighbor is in the `Established` stage. If the neighbor has never moved into the `Established` stage, the word never displays. The output format is:

  ```
  1 day = 00:12:23 (hours:minutes:seconds), 1 week = 1d21h (DaysHours), 1 week
  + 11w2d (WeeksDays)
  ```

- `State/Pfxrcd`—If the neighbor is in the `Established` stage, this is the number of network prefixes received. If you configured a maximum limit using the `neighbor maximum-prefix` command, `prfxd` appears in this column. If the neighbor is not in the `Established` stage, the current stage - `Idle`, `Connect`, `Active`, `OpenSent`, `OpenConfirm` displays. When the peer is transitioning between states and clearing the routes received, the phrase `Purging` may appear in this column. If the neighbor is disabled, the phrase `Admin shut` appears in this column.

The suppressed status of aggregate routes may not display in the command output.

Example

```
OS10# show ip bgp summary
BGP router identifier 80.1.1.1 local AS number 102
Neighbor    AS    MsgRcvd  MsgSent  Up/Down    State/Pfx
80.1.1.2    800   24       23       00:09:15   5
```

Supported Releases 10.2.0E or later

show ip route

Displays information about IPv4 BGP routing table entries.

Syntax `show ip route`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information This command displays information about IPv4 BGP routing table entries.

Example

```
OS10# show ip route
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
```

```
       E2 - OSPF external type 2, * - candidate default,
       + - summary route, > - non-active route
Gateway of last resort is not set
  Destination                Gateway                                  Dist/Metric       L
-------------------------------------------------------------------------------------------
  C     10.1.1.0/24         via 10.1.1.1         ethernet1/1/17      0/0
  B  IN 100.1.1.0/24        via 10.1.1.2                             200/0
  B  IN 101.1.1.0/24        via 10.1.1.2                             200/0
  B  IN 102.1.1.0/24        via 10.1.1.2                             200/0
  B  IN 103.1.1.0/24        via 10.1.1.2                             200/0
  B  IN 104.1.1.0/24        via 10.1.1.2                             200/0
```

Supported Releases 10.4.2.0 or later

show ipv6 route

Displays information about IPv6 BGP routing table entries.

Syntax `show ipv6 route`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information This command displays information about IPv6 BGP routing table entries.

Example

```
OS10# show ipv6 route
```

Supported Releases 10.4.2.0 or later

soft-reconfiguration inbound

Enables soft-reconfiguration for a neighbor.

Syntax `soft-reconfiguration inbound`

Parameters None

Default Not configured

Command Modes ROUTER-BGP-NEIGHBOR-AF

Usage Information This command is not supported on a peer-group level. To enable soft-reconfiguration for peers in a peer-group, you must enable this command at a per-peer level. With soft-reconfiguration inbound, all updates received from this neighbor are stored unmodified, regardless of the inbound policy. When inbound soft-reconfiguration is performed later, the stored information generates a new set of inbound updates. The `no` version of this command disables soft-reconfiguration inbound for a BGP neighbor.

Example (IPv4)

```
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-bgp-neighbor-af)# soft-reconfiguration inbound
```

Example (IPv6)

```
OS10(conf-router-neighbor)# address-family ipv6 unicast
OS10(conf-router-bgp-neighbor-af)# soft-reconfiguration inbound
```

Supported Releases 10.3.0E or later

template

Creates a peer-group template to assign it to BGP neighbors.

Syntax `template template-name`

Parameters *template-name* — Enter a peer-group template name. A maximum of 16 characters.

Default Not configured

Command Mode CONFIG-ROUTER-BGP

Usage Information Members of a peer-group template inherit the configuration properties of the template and share the same update policy. The `no` version of this command removes a peer-template configuration.

NOTE: To configure these settings for a non default VRF instance, you must first enter the ROUTER-CONFIG-VRF sub mode using the following commands:

1. Enter the ROUTER BGP mode using the `router bgp as-number` command.
2. From the ROUTER BGP mode, enter the ROUTER BGP VRF mode using the `vrf vrf-name` command.

Example

```
OS10(conf-router-bgp-10)# template solar
OS10(conf-router-bgp-template)#
```

Supported Releases 10.3.0E or later

timers

Adjusts BGP keepalive and holdtime timers.

Syntax `timers keepalive holdtime`

Parameters

- *keepalive*—Enter the time interval, in seconds, between keepalive messages sent to the neighbor routers, from 1 to 65535.
- *holdtime*—Enter the time interval, in seconds, between the last keepalive message and declaring a router dead, from 3 to 65535.

Default `keepalive` 60 seconds; `holdtime` 180 seconds

Command Mode ROUTER-BGP

Usage Information The configured timer value becomes effective after a BGP hard reset. The timer values negotiate from peers. The `no` version of this command resets the value to the default.

Example

```
OS10(conf-router-bgp)# timers 30 90
```

Supported Releases 10.3.0E or later

vrf

Enters the CONFIG-ROUTER-VRF command mode.

Syntax `vrf vrf-name`

Parameters	None
Default	None
Command Mode	ROUTER-BGP
Usage Information	This mode allows you to apply BGP configurations to non-default VRFs.

Example	`OS10(config)#router bgp 100`
	`OS10(config-router-bgp-100)#`
	`OS10(config-router-bgp-100)#vrf vrf_test1`
	
	`OS10(config-router-bgp-100-vrf)#`

Supported Releases	10.3.0E or later

weight

Assigns a default weight for routes from the neighbor interfaces.

Syntax	`weight` *`number`*
Parameters	*`number`*—Enter a number as the weight for routes, from 1 to 4294967295.
Default	0
Command Mode	ROUTER-BGP-NEIGHBOR
Usage Information	The path with the highest weight value is preferred in the best-path selection process. The `no` version of this command resets the value to the default.
Example	`OS10(conf-router-bgp-neighbor)# weight 4096`
Supported Releases	10.3.0E or later

Equal cost multi-path

ECMP is a routing technique where next-hop packet forwarding to a single destination occurs over multiple best paths. When you enable ECMP, OS10 uses a hash algorithm to determine the next-hop. The hash algorithm makes hashing decisions based on values in various packet fields and internal values.

Configure the hash algorithm in CONFIGURATION mode.

```
hash-algorithm ecmp crc
```

NOTE: The S5148F-ON platform supports only the crc hash algorithm.

Change hash algorithm

```
OS10(config)# hash-algorithm ecmp crc
```

Load balancing

To increase bandwidth, traffic is balanced across member links. RTAG7 is a hash algorithm that load balances traffic within a trunk group in a controlled manner. RTAG7 balances traffic to more effectively use member links as traffic gets more diverse.

RTAG7 generates a hash that consists of two parts:

- The first part generates from packet headers to identify micro-flows in traffic. By default, all listed parameters are enabled for load balancing except the ingress port.

```
OS10# show load-balance

Load-Balancing Configuration For LAG and ECMP:
-----------------------------------------------
IPV4 Load Balancing          : Enabled
IPV6 Load Balancing          : Enabled
MAC Load Balancing           : Enabled
TCP-UDP Load Balancing       : Enabled
Ingress Port Load Balancing : Disabled
IPV4 FIELDS    : source-ip destination-ip protocol vlan-id l4-destination-port l4-source-port
IPV6 FIELDS    : source-ip destination-ip protocol vlan-id l4-destination-port l4-source-port
MAC FIELDS     : source-mac destination-mac  ethertype  vlan-id
TCP-UDP FIELDS: l4-destination-port  l4-source-port
```

- The second part generates from the static physical configuration such as the ingress and egress port numbers.

To generate load balancing based on any parameters, change the hash field using the `load-balance` command.

Resilient hashing

To increase bandwidth and for load balancing, traffic is distributed across the next hops of an ECMP group or member ports of a port channel. OS10 uses a hash algorithm to determine a hash key. The egress port in a port channel or the next hop in an ECMP group is selected based on the hash key modulo the number of ports in a port channel or next hops in an ECMP group, respectively. When a member link goes down or a new member link gets added, the traffic flows are remapped based on the new hash result.

In this section, the term, "member link" refers to either a member physical port, in the case of port channels or next hop in the case of ECMP groups.

With resilient hashing, when a member link goes down, the existing flows are not affected (remapped). Resilient hashing reassigns the traffic from the failed link to another member link without remapping the other existing flows. However, minimal re-mapping occurs when a new member link is added.

Resilient hashing is supported both for Port Channels and Equal Cost MultiPath Groups. Resilient hashing is a global configuration and can be configured for both port channels and ECMP independently.

NOTE:

- Resilient hashing is supported only on certain platforms.
- The flow-map table always has an even number of entries.

To enable resilient hashing for ECMP groups, enter the following command in CONFIGURATION mode:

```
OS10(config)# enhanced-hashing resilient-hashing ecmp
```

```
OS10(config)# enhanced-hashing resilient-hashing lag
```

Supported platforms

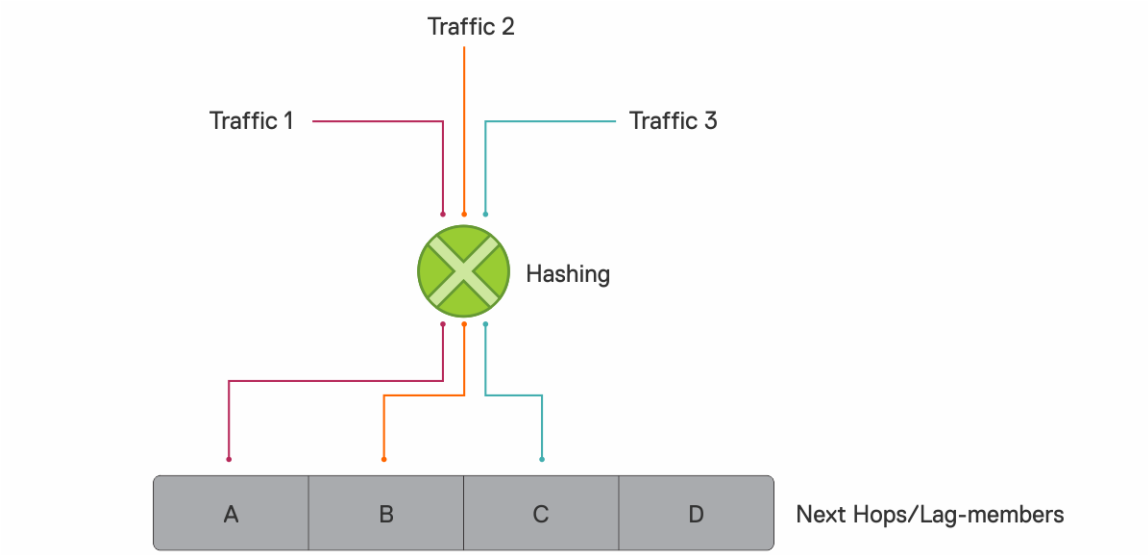
The following table lists the platforms that support the resilient hashing feature.

Table 15. Supported platforms for resilient hashing feature

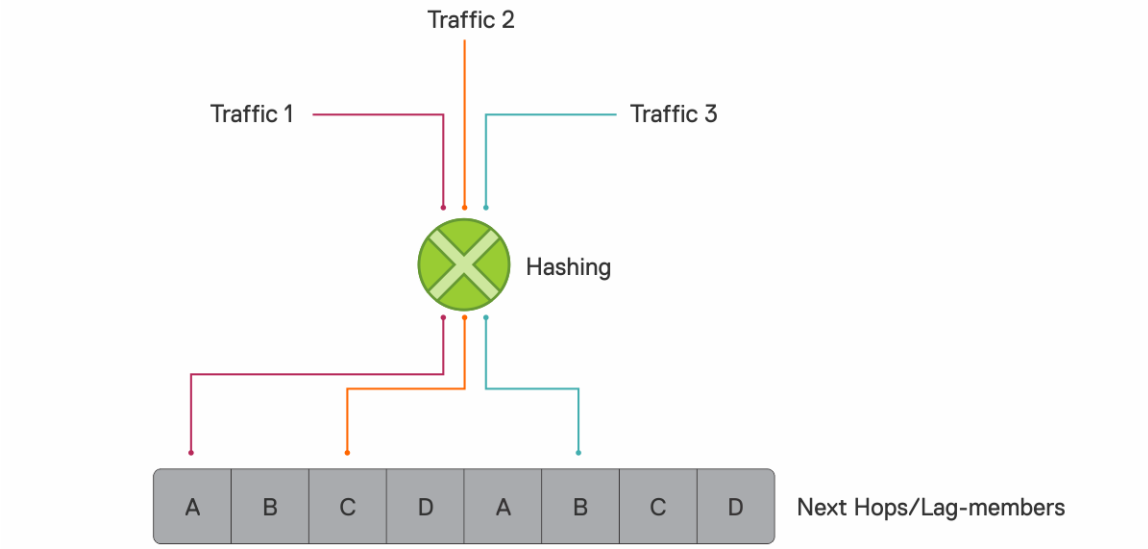
Platform	Resilient hashing on ECMP	Resilient hashing on port channels
S6000-ON, S6010-ON, S4048T-ON, S4100-ON Series, S5200F-ON Series	Y	Y
Z9100-ON Series, Z9200-ON Series	Y	N
Mx Series	Y	N
S3000-ON, S4200-ON Series, S5148F-ON	N	N

Examples

Normal traffic flow without resilient hashing

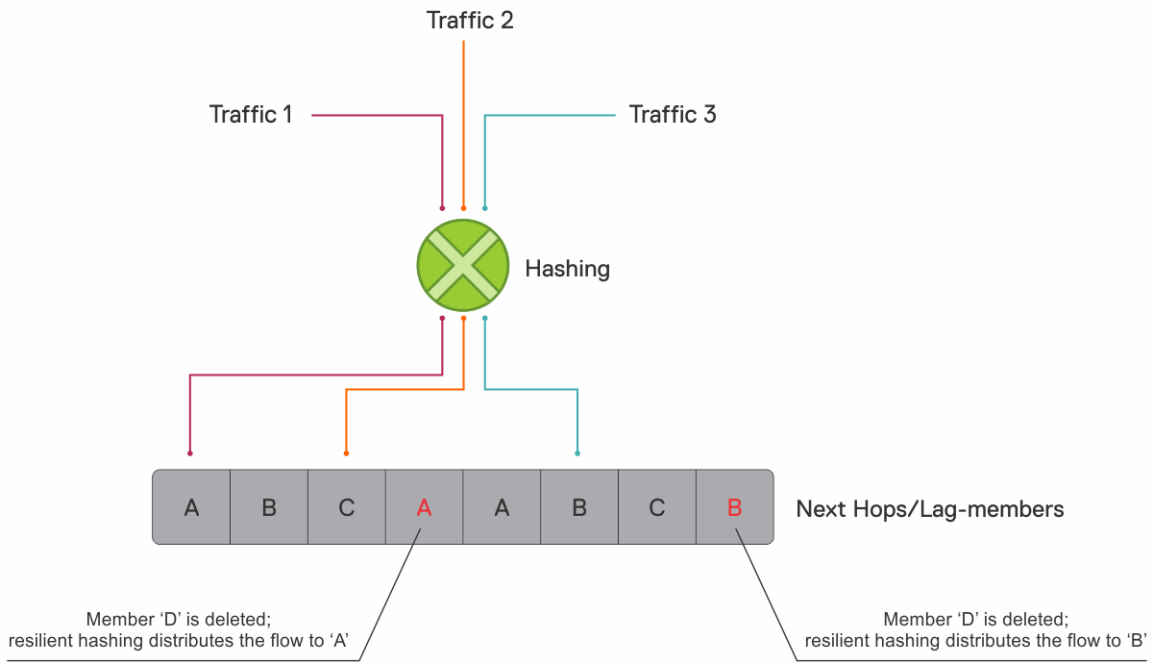


Traffic flow with resilient hashing enabled



When you enable resilient hashing for ECMP groups, the flow-map table gets created with 64 paths (which is the OS10 default maximum number of ECMP paths) and traffic is equally distributed. In the above example, we see traffic 1 mapped to next hop 'A'; traffic 2 mapped to next hop 'C'; and traffic 3 mapped to next hop 'B.'

Member link goes down

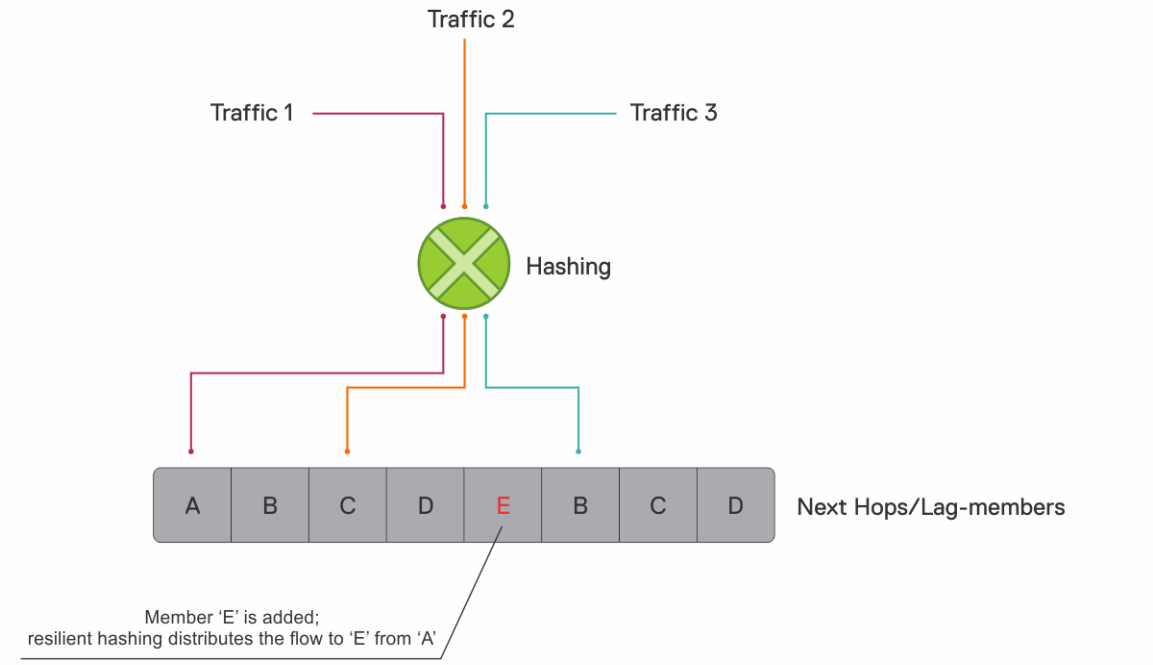


In this example, when member 'D' goes down, resilient hashing distributes the traffic intended for member link 'D' to A and B. The existing traffic, traffic 1, 2, and 3 are not disturbed.

Member link is added

However, when a new member is added, resilient hashing does some minimal remapping for better load balancing as shown in the following image.

Adding a member



Important notes

- Resilient hashing on port channels is applicable only for unicast traffic.
- For resilient hashing on ECMP groups, the ECMP path should be in multiples of 64. Before you enable resilient hashing, ensure that the maximum ECMP path is set to a multiple of 64. You can configure this value using the ip ecmp-group maximum-paths command.

Maximum ECMP groups and paths

The maximum number of ECMP groups supported on the switch depends on the maximum ECMP paths configured on the switch. To view the maximum number of ECMP groups and paths, use the `show ip ecmp-group details` command.

```
OS10# show ip ecmp-group details
Maximum Number of ECMP Groups : 256
Maximum ECMP Path per Group : 64
Next boot configured Maximum ECMP Path per Group : 64
```

The default value for the maximum number of ECMP paths per group is 64. This value is configurable and you can configure a maximum of up to 128 ECMP paths per group.

The `Maximum ECMP Path per Group` is the current value configured in the hardware. The `Next boot configured Maximum ECMP Path per Group` is the value that is configured for maximum ECMP path and will take effect after the next reboot.

You can increase or decrease the maximum number of ECMP groups using the `ip ecmp-group maximum-paths` *number* command. The number of ECMP groups is inversely proportional to the number of ECMP paths.

To configure maximum paths per ECMP route:

```
OS10# configure terminal
OS10(config)# ip ecmp-group maximum-paths 10
OS10(config)# exit
```

```
OS10# write memory
OS10# reload
```

ECMP commands

enhanced-hashing

Ensures that existing traffic flows are not re-mapped when a member link goes down.

Syntax `enhanced-hashing resilient-hashing {lag | ecmp}`

Parameters

- `resilient-hashing`—Enter the keyword `resilient-hashing`.
- `{ecmp | lag}`—Enter the keyword `ecmp` or `lag` to enable resilient hashing for port channel or ECMP group.

Defaults Disabled

Command Mode CONFIGURATION

Usage Information The `no` version of this command disables resilient hashing.

Example

```
OS10(config)# enhanced-hashing resilient-hashing ecmp
OS10(config)# enhanced-hashing resilient-hashing lag
```

Supported Releases 10.4.3.0 or later

Related link

Supported platforms

hash-algorithm

Changes the hash algorithm that distributes traffic flows across ECMP paths and the link aggregation group (LAG).

Syntax `hash-algorithm {ecmp | lag | seed {seed-value}} {crc | crc16cc | crc32LSB | crc32MSB | xor | xor1 | xor2 | xor4 | xor8 | random}`

Parameters

NOTE: The S5148F-ON platform supports only the crc parameter.

- `ecmp`—Enables the ECMP hash configuration.
- `lag`—Enables the LAG hash configuration for Layer 2 (L2) only.
- `seed`—Changes the hash algorithm seed value to get a better hash value.
- *`seed-value`*—Enter a hash algorithm seed value, from 0 to 4294967295.
- `crc`—Enables the cyclic redundancy check (CRC) polynomial for hash computation.
- `crc16cc`—16 bit CRC16 using CRC16-CCITT polynomial
- `crc32LSB`—LSB 16 bits of computed CRC32(default)
- `crc32MSB`—MSB 16 bits of computed CRC32
- `xor` — Enables upper 8 bits of CRC and lower 8 bits of XOR value for computation.
- `xor1`—Enables upper 8 bits of CRC16-BISYNC and lower 8 bits of xor1

- `xor2`—Enables upper 8 bits of CRC16-BISYNC and lower 8 bits of xor2
- `xor4`—Enables upper 8 bits of CRC16-BISYNC and lower 8 bits of xor4
- `xor8`—Enables upper 8 bits of CRC16-BISYNC and lower 8 bits of xor8
- `random` — Enables a hash algorithm random seed value for ECMP or LAG hash computation.

Default `crc`

Command Mode CONFIGURATION

Usage Information The hash value calculated with this command is unique to the entire system. Different hash algorithms are based on the number of port-channel members and packet values. The default hash algorithm yields the most balanced results in various test scenarios, but if the default algorithm does not provide a satisfactory distribution of traffic, use this command to designate another algorithm.

When a port-channel member leaves or is added to the port-channel, the hash algorithm recalculates to balance traffic across the members. The `no` version of this command returns the value to the default.

Example

```
OS10(config)# hash-algorithm lag crc
```

Supported Releases 10.3.0E or later

ip ecmp-group maximum-paths

Configures the maximum number of ECMP paths per route.

Syntax `ip ecmp-group maximum-paths` *`number`*

Parameters *`number`* — Enter the maximum number of ECMP paths, from 2 to 128.

Default 64

Command Mode CONFIGURATION

Usage Information To save the new ECMP settings, use the `write memory` command, then reload the system for the new settings to take effect. The `no` version of this command returns the value to the default.

Example

```
OS10# configure terminal
OS10(config)# ip ecmp-group maximum-paths 2
OS10(config)# exit
OS10# write memory
OS10# reload
```

Supported Releases 10.4.3.0 or later

link-bundle-utilization trigger-threshold

Configures a threshold value to trigger traffic monitoring distribution on an ECMP link bundle.

Syntax `link-bundle-utilization trigger-threshold` *`value`*

Parameters *`value`* — Enter a link bundle trigger threshold value, from 0 to 100.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command disables the configuration.

Example

```
OS10(config)# link-bundle-utilization trigger-threshold 80
```

Supported Releases 10.2.0E or later

load-balancing

Distributes or load balances incoming traffic using the default parameters in the hash algorithm.

Syntax

```
load-balancing {ingress-port enable | [tcp-udp-selection l4-destination-port |
l4-source-port] | [ip-selection destination-ip | source-ip | protocol | vlan-id
| l4-destination-port | l4-source-port] | [ipv6-selection destination-ip |
source-ip | protocol | vlan-id | l4-destination-port | l4-source-port] | [mac-
selection destination-mac | source-mac | ethertype | vlan-id]}
```

Parameters

- `ingress-port enable` — Enables load-balancing on ingress ports.
- `tcp-udp-selection` — Enables the TCP UDP port for the load-balancing configuration.
- `ip-selection` — Enables IPv4 key parameters to use in the hash computation.
- `ipv6-selection` — Enables IPV6 key parameters to use in hash computation.
- `destination-ip` — Enables the destination IP address in the hash calculation.
- `source-ip` — Enables the source IP address in the hash calculation.
- `protocol` — Enables protocol information in the hash calculation.
- `vlan-id` — Enables VLAN ID information in the hash calculation.
- `l4-destination-port` — Enables Layer 4 (L4) destination port information in the hash calculation.
- `l4-source-port` — Enables L4 source port information in the hash calculation.
- `mac-selection` — Enables MAC load-balancing configurations.
- `destination-mac`— Enables destination MAC information in the hash calculation.
- `source-mac`— Enables source MAC information in the hash calculation.
- `ethertype` — Enables Ethernet type information in the hash calculation.

Default

- `ip-selelection—source—ip dest-ip vlan-id l4-source-port l4-dest-port ipv4 protocol`
- `ipv6-selection—source—ipv6 dest-ipv6 vlan—id l4-source-port l4-dest-port ipv6 protocol`
- `mac-selection—source—mac destination—mac vlan-id ethertype`
- `tcp-udp-selection—l4-source-port l4-dest-port`

Command Mode CONFIGURATION

Usage Information

- IPv4– selection: `source-ip destination-ip protocol vlan-id l4-destination-port l4-source-port`
- IPv6 destination address: `source-ip destination-ip protocol vlan-id l4-destination-port l4-source-port`
- MAC parameters: `source-mac destination-mac ethertype vlan-id`
- TCP/UDP parameters: `l4-destination-port l4-source-port`

The `no` version of this command resets the value to the default.

Example (Ingress)

```
OS10(config)# load-balancing ingress-port enable
```

Example (IP Selection)	`OS10(config)# load-balancing ip-selection destination-ip source-ip`
Supported Releases	10.2.0E or later

show enhanced-hashing resilient-hashing

Displays the status of `enhanced-hashing` command.

Syntax	`show enhanced-hashing resilient-hashing {lag | ecmp}`
Parameters	None
Default	Disabled
Command Mode	EXEC
Usage Information	None

Example	`OS10# show enhanced-hashing resilient-hashing lag`
	`Resilient Hashing Configuration For LAG:`
	`----------`
	`LAG Resilient hashing : Disabled`
	
	`OS10# show enhanced-hashing resilient-hashing ECMP`
	`Resilient Hashing Configuration For ECMP:`
	`----------`
	`ECMP Resilient hashing : Disabled`

Supported Releases	10.4.3.0 or later

show hash-algorithm

Displays hash-algorithm information.

Syntax	`show hash-algorithm`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	None
Example	`OS10# show hash-algorithm` `EcmpAlgo - crc LabAlgo - crc`
Supported Releases	10.3.0E or later

show ip ecmp-group details

Displays the number of ECMP groups and paths.

Syntax	`show ip ecmp-group details`
Parameters	None
Default	Not configured

Command Mode	EXEC
Usage Information	None

Example	`OS10# show ip ecmp-group details`
	`Maximum Number of ECMP Groups : 256`
	`Maximum ECMP Path per Group : 64`
	`Next boot configured Maximum ECMP Path per Group : 64`

Supported Releases	10.4.3.0 or later

show load-balance

Displays the global traffic load-balance configuration.

Syntax	`show load-balance`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	None

Example	`OS10# show load-balance`
	
	`Load-Balancing Configuration For LAG & ECMP:`
	`-------------------------------------------`
	`IPV4 Load Balancing Enabled`
	`IPV4 FIELDS : source-ipv4 dest-ipv4 vlan protocol L4-source-port L4-dest-port`
	
	`IPV6 Load Balancing Enabled`
	`IPV6 FIELDS : source-ipv6 dest-ipv6 vlan protocol L4-source-port L4-dest-port`
	
	`Mac Load Balancing Enabled`
	`MAC FIELDS : source-mac dest-mac vlan ethertype`
	
	`mac-in-mac header based hashing is disabled`
	`TcpUdp Load Balancing Enabled`

Supported Releases	10.3.0E or later

IPv4 routing

OS10 supports IPv4 addressing including variable-length subnetting mask (VLSM), Address Resolution Protocol (ARP), static routing, and routing protocols. With VLSM, you can configure one network with different masks. You can also use supernetting, which increases the number of subnets. You can add a mask to the IP address to separate the network and host portions of the IP address to add a subnet.

You need to configure IPv4 routing for IP hosts to communicate with one another in the same network, or in different networks.

Assign interface IP address

You can assign primary and secondary IP addresses to a physical or logical interface to enable IP communication between the system and hosts connected to a specific interface. Assign one primary address and secondary IP addresses to each interface. By default, all ports are in the default VLAN—VLAN 1.

1 Enter the interface type information to assign an IP address in CONFIGURATION mode.

```
interface interface
```

- `ethernet`—Physical interface
- `port-channel`—Port-channel ID number
- `vlan`—VLAN ID number
- `loopback`—Loopback interface ID
- `mgmt`—Management interface

2 Enable the interface in INTERFACE mode.

```
no shutdown
```

3 Remove the interface from the default VLAN in INTERFACE mode.

```
no switchport
```

4 Configure a primary IP address and mask on the interface in INTERFACE mode.

```
ip address ip-address mask [secondary]
```

- *ip-address mask*—Enter the IP address in dotted decimal format—A.B.C.D. and mask in slash prefix-length format (/24).
- `secondary`—Enter a secondary backup IP address for the interface.

Assign interface IP address to interface

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# ip address 10.10.1.4/8
```

View interface configuration

```
OS10# show interface ethernet 1/1/1
Ethernet 1/1/1 is up, line protocol is up
Hardware is Dell EMC Eth, address is 00:0c:29:98:1b:79
    Current address is 00:0c:29:98:1b:79
Pluggable media present, QSFP+ type is QSFP+ 40GBASE CR 1.0M
    Wavelength is 64
    SFP receive power reading is 0.0
Interface index is 16866084
Internet address is not set
Mode of IPv4 Address Assignment: not set
MTU 1532 bytes
LineSpeed 40G, Auto-Negotiation on
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout: 60
Last clearing of "show interface" counters: 3 weeks 1 day 23:12:50
Queuing strategy: fifo
Input statistics:
     0 packets, 0 octets
     0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     0 Multicasts, 0 Broadcasts, 0 Unicasts
     0 runts, 0 giants, 0 throttles
     0 CRC, 0 overrun, 0 discarded
Output statistics:
     0 packets, 0 octets
     0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     0 Multicasts, 0 Broadcasts, 0 Unicasts
     0 throttles, 0 discarded, 0 Collisions, 0 wreddrops
Rate Info(interval 299 seconds):
     Input 0 Mbits/sec, 0 packets/sec, 0% of line rate
     Output 0 Mbits/sec, 0 packets/sec, 0% of line rate
Time since last interface status change: 3 weeks 1 day 20:54:37
```

Configure static routing

You can configure a manual or static route for open shortest path first (OSPF).

- Configure a static route in CONFIGURATION mode.
  ```
  ip route ip-prefix/mask {next-hop | interface interface [route-preference]}
  ```
  - *ip-prefix*—IPv4 address in dotted decimal in A.B.C.D format.
  - *mask*—Mask in slash prefix-length format (/X).
  - *next-hop*—Next-hop IP address in dotted decimal in A.B.C.D format.
  - *interface*—Interface type with the node/slot/portinformation
  - *route-preference*—(Optional) Route-preference range, from 1 to 255.

Configure static routes

```
OS10(config)# ip route 200.200.200.0/24 10.1.1.2
```

View configured static routes

```
OS10# show ip route static
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF,IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, > - non-active route
Gateway of last resort is not set
Destination        Gateway                  Dist/Metric Last Change
-----------------------------------------------------------------
S  200.200.200.0/24 via 10.1.1.2 ethernet1/1/1  0/0      00:00:03
```

OS10 installs a static route if the next hop is on a directly connected subnet. A next-hop that is not on the directly connected subnet which recursively resolves to a next-hop on the interface's configured subnet also automatically configures. For example, if `interface ethernet 1/1/5` has IP address on subnet 100.0.0.0/8, and if 10.1.1.0/24 recursively resolves to 100.1.1.1, the system installs the static route:

- When the interface goes down, OS10 withdraws the route.
- When the interface comes up, OS10 reinstalls the route.
- When the recursive resolution is *broken*, OS10 withdraws the route.
- When the recursive resolution is satisfied, OS10 reinstalls the route.

Address Resolution Protocol

Address Resolution Protocol (ARP) runs over Ethernet and enables end stations to learn the MAC addresses of neighbors on an IP network. Using ARP, OS10 automatically updates the *ARP cache* table that maps the MAC addresses to their corresponding IP addresses. The *ARP cache* enables dynamically learned addresses to be removed after a time period you configure.

Configure static ARP entries

You can manually configure static entries in the ARP mapping table. Dynamic ARP is vulnerable to spoofing. To avoid spoofing, configure static entries. Static entries take precedence over dynamic ARP entries.

1 Configure an IP address and MAC address mapping for an interface in INTERFACE mode.
  ```
  ip arp ip-address mac address
  ```

- *ip-address*—IP address in dotted decimal format in A.B.C.D format.
- *mac address*—MAC address in nnnn.nnnn.nnnn format

These entries do not age, and you can only remove them manually. To remove a static ARP entry, use the `no arp ip-address` command.

Configure static ARP entries

```
OS10(config)# interface ethernet 1/1/6
OS10(conf-if-eth1/1/6)# ip arp 10.1.1.5 08:00:20:b7:bd:32
```

View ARP entries

```
OS10# show ip arp interface ethernet 1/1/6
Protocol    Address    Hardware Interface   Interface       VLAN
-------------------------------------------------------------------
Internet    10.1.1.5   08:00:20:b7:bd:32    ethernet1/1/6
```

IPv4 routing commands

clear ip arp

Clears the dynamic ARP entries from a specific interface or optionally delete (no-refresh) ARP entries from the content addressable memory (CAM).

Syntax `clear ip arp [vrf vrf-name] [interface interface | ip ip-address] [no-refresh]`

Parameters

- `vrf vrf-name` — (Optional) Enter `vrf` then the name of the VRF to clear ARP entries corresponding to that VRF.
- `interface interface`— (Optional) Specify an interface type:
  - `ethernet` — Physical interface.
  - `port-channel` — Port-channel identifier.
  - `vlan` — VLAN identifier.
  - `loopback` — Loopback interface identifier.
  - `virtual-network vn-id` — Virtual network ID.
- `ip ip-address` — (Optional) Specify the IP address of the ARP entry to clear.
- `no-refresh` — (Optional) Delete the ARP entry from CAM. You can also use this option with *interface* or `ip` *ip-address* to specify which dynamic ARP entries to delete.

Default Not configured

Command Mode EXEC

Usage Information Transit traffic may not forward during the period when deleted ARP entries resolve again and re-install in CAM.

NOTE: Use this option with extreme caution.

Example

```
OS10# clear ip arp interface ethernet 1/1/5
```

Supported Releases 10.2.0E or later

clear ip route

Clears the specified routes from the IP routing table.

Syntax

```
clear ip route [vrf vrf-name] {* | A.B.C.D/mask}
```

Parameters

- `vrf vrf-name` — (Optional) Enter the keyword `vrf` and then the name of the VRF to clear the routes corresponding to that VRF.
- `*`—Clear the entire IP routing table. This option refreshes all the routes in the routing table. Traffic flow is affected for all the routes in the switch.
- *A.B.C.D/mask* —Specify the IP route to remove from the IP routing table. This option refreshes all the routes in the routing table. Traffic flow is affected only for the specified route in the switch.

Default Not configured

Command Mode EXEC

Usage Information This command does not remove the static routes from the routing table.

Example

```
OS10# clear ipv6  route 10.1.1.0/24
```

Supported Releases 10.3.0E or later

ip address

Configure the IP address to an interface.

Syntax

```
ip address ip-address/mask
```

Parameters *ip–address/mask* — Enter the IP address.

Defaults None

Command Mode INTERFACE

Usage Information The `no` version of this command removes the IP address set for the interface.

Example

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ip address 10.1.1.0/24
```

Supported Releases 10.3.0E or later

ip address dhcp

Enables DHCP client operations on the interface.

Syntax

```
ip address dhcp
```

Parameters None

Defaults None

Command Mode INTERFACE

Usage Information The `no` version of this command disables DHCP operations on the interface.

Example

```
OS10(config)# interface mgmt 1/1/1
OS10(conf-if-ma-1/1/1)# ip address dhcp
```

Supported Releases 10.3.0E or later

ip arp

Configures static ARP and maps the IP address of the neighbor to a MAC address.

Syntax `ip arp mac-address`

Parameters *mac-address* — Enter the MAC address of the IP neighbor in A.B.C.D format.

Default Not configured

Command Mode INTERFACE

Usage Information Do not use Class D (multicast) or Class E (reserved) IP addresses. Zero MAC addresses (00:00:00:00:00:00) are invalid. The `no` version of this command disables the IP ARP configuration.

Example

```
OS10(conf-if-eth1/1/6)# ip arp 10.1.1.5 08:00:20:b7:bd:32
```

Supported Releases 10.2.0E or later

ip arp gratuitous

Enables an interface to receive or send gratuitous ARP requests and updates.

Syntax `ip arp gratuitous {update | request}`

Parameters

- `update` — Specify to enable or disable ARP cache updates for gratuitous ARP.
- `request` — Specify to enable or disable sending gratuitous ARP requests when duplicate address is detected.

Default Not configured

Command Mode CONFIG-INTERFACE

Usage Information When a reply to a gratuitous ARP request is received, it indicates an IP address conflict in the network. The `no` version of this command disables the ARP cache updates for gratuitous ARP.

Example

```
OS10(conf-if-eth1/1/6)# ip arp gratuitous update
OS10(conf-if-eth1/1/6)# ip arp gratuitous request
```

Supported Releases 10.2.0E or later

ip route

Assigns a static route on the network device.

Syntax `ip route [dst-vrf vrf-name] ip-prefix mask {next-hop | interface interface-type [route-preference]}`

Parameters

- `dst-vrf vrf-name` — (Optional) Enter `vrf` and then the name of the VRF to configure a static route corresponding to that VRF. Use this VRF option after the `ip route` keyword to configure a static route on that specific VRF.

- *ip-prefix* — Enter the IP prefix in dotted decimal A.B.C.D format.
- *mask* — Enter the mask in slash prefix-length /x format.
- *next-hop* — Enter the next-hop IP address in dotted decimal A.B.C.D format.
- `interface` *interface-type* — Enter the interface type and interface information. The interface types supported are: Ethernet, port-channel, VLAN, and Null.
- *route-preference* — (Optional) Enter the range, from 1 to 255.

Default Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command deletes a static route configuration.

Example

```
OS10(config)# ip route 200.200.200.0/24 10.1.1.2
```

```
OS10(config)# ip route 200.200.200.0/24 interface null 0
```

Supported Releases 10.2.0E or later

show ip arp

Displays the ARP table entries for a specific IP address or MAC address, static, dynamic, and a summary of all ARP entries.

Syntax

```
show ip arp [vrf vrf-name] [interface [ethernet | vlan | port-channel] | ip-address | mac-address | static | dynamic | summary]
```

Parameters

- `vrf` *vrf-name* — Enter `vrf` then the name of the VRF to display ARP entries corresponding to that VRF.
- `interface` — (Optional) Enter the keyword and interface information:
  - *ethernet* — Enter the node/slot/port[:subport] information.
  - *vlan* — Enter the VLAN ID number, from 1 to 4093.
  - *port-channel* — Enter the port-channel ID number, from 1 to 128.
- *ip-address* — (Optional) Enter the IP address for the ARP entry in A.B.C.D format.
- *mac-address* — (Optional) Enter the MAC address in nn:nn:nn:nn:nn:nn format.
- `static` — (Optional) Enter the keyword to display static ARP entries.
- `dynamic` — (Optional) Enter the keyword to display dynamic ARP entries.
- `summary` — (Optional) Enter the keyword to display a summary of all ARP entries.

Default Not configured

Command Mode EXEC

Usage Information This command shows both static and dynamic ARP entries.

Example (IP Address)

```
OS10# show ip arp 192.168.2.2
Protocol    Address       Hardware Interface   Interface       VLAN
------------------------------------------------------------------
Internet    192.168.2.2   00:01:e8:8b:3c:01    ethernet1/1/6
```

Example (Static)

```
OS10# show ip arp summary
Total Entries   Static Entries    Dynamic Entries
-------------------------------------------------
     3994               0                3994
OS10# show ip arp 100.1.2.1
Protocol       Address        Hardware Interface    Interface       VLAN
-------------------------------------------------------------------------
Internet       100.1.2.1      00:a0:c9:00:01:04     port-channel11  1002
```

```
OS10# show ip arp dynamic
Protocol       Address         Hardware Interface    Interface       VLAN
-------------------------------------------------------------------------
Internet       9.0.0.2         00:24:00:00:00:00     ethernet1/1/10  4001
Internet       100.1.1.1       00:a0:c9:00:01:04     port-channel11  1001
Internet       100.1.2.1       00:a0:c9:00:01:04     port-channel11  1002
Internet       100.1.3.1       00:a0:c9:00:01:04     port-channel11  1003
Internet       100.1.4.1       00:a0:c9:00:01:04     port-channel11  1004
Internet       100.1.5.1       00:a0:c9:00:01:04     port-channel11  1005
Internet       100.1.6.1       00:a0:c9:00:01:04     port-channel11  1006
Internet       100.1.7.1       00:a0:c9:00:01:04     port-channel11  1007
Internet       100.1.8.1       00:a0:c9:00:01:04     port-channel11  1008
```

Example (Dynamic)

```
OS10# show ip arp dynamic
Protocol    Address          Hardware Interface   Interface       VLAN
-------------------------------------------------------------------------
Internet    10.16.127.143    00:01:e8:75:c1:bb    ethernet1/1/6   -
Internet    192.168.1.2      00:01:e8:8b:3c:01    ethernet1/1/1   100
```

Supported Releases 10.2.0E or later

show ip route

Displays IP route information.

Syntax

```
show ip route [vrf vrf-name] [all | bgp | connected | ospf process-id  | static | ip-pr
summary]
```

Parameters

- `vrf vrf-name` — (Optional) Enter `vrf` and then the VRF name to list the routes in the route table of a specific VRF.
- `all` — (Optional) Displays both active and non-active IP routes.
- `bgp` — (Optional) Displays BGP route information.
- `connected` — (Optional) Displays only the directly connected routes.
- `ospf process-id` — (Optional) Displays route information for the OSPF process, from 1 to 65535.
- `static` — (Optional) Displays static route information.
- *ip-prefix/mask* — (Optional) Displays routes for the destination prefix list.

  NOTE: This option works only for the exact prefix and the mask length.
- `summary` — (Optional) Displays an IP route summary.

Defaults Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip route
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, > - non-active route
Gateway of last resort is not set
 Destination         Gateway             Dist/Metric  Last Change
-------------------------------------------------------------------
 C     10.1.1.0/24   via 10.1.1.1 vlan100   0/0        01:16:56
 B EX  10.1.2.0/24   via 10.1.2.1 vlan101   20/0       01:16:56
```

```
 O      10.1.3.0/24    via 10.1.3.1 vlan102    110/2       01:16:56
 B IN   10.1.4.0/24    via 10.1.4.1 vlan103    200/0       01:16:56
```

```
OS10(config)# do show ip route vrf VRF1
Codes: C - connected
S - static
B - BGP, IN - internal BGP, EX - external BGP
O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
E2 - OSPF external type 2, * - candidate default,
+ - summary route, > - non-active route
Gateway of last resort is not set
Destination Gateway Dist/Metric Last Change
--------------------------------------------------------------------------------
C 120.0.0.0/24 via 120.0.0.1 ethernet1/1/1 0/0 00:00:57
S 160.0.0.0/24 via 120.0.0.2 ethernet1/1/1 1/0 00:00:04
OS10(config)# do show ip route vrf VRF2
Codes: C - connected
S - static
B - BGP, IN - internal BGP, EX - external BGP
O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
E2 - OSPF external type 2, * - candidate default,
+ - summary route, > - non-active route
Gateway of last resort is not set
Destination Gateway Dist/Metric Last Change
--------------------------------------------------------------------------------
C 140.0.0.0/24 via 140.0.0.1 ethernet1/1/2 0/0 00:01:54
B  IN 160.0.0.0/24   via 120.0.0.2   200/0             00:00:02
```

Supported Releases 10.2.0E or later

IPv6 routing

OS10 supports IPv6 routing and addressing, including the Neighbor Discovery Protocol (NDP), stateless IPv6 address autoconfiguration, and stateful IPv6 address configuration. Configure IPv6 routing for IP hosts to communicate with one another in the same network, or in different networks.

Enable or disable IPv6

By default:

- IPv6 forwarding is enabled on physical Ethernet interfaces, VLANs, and port groups. IPv6 forwarding is disabled only when you enable IPv6 address autoconfiguration on an interface and set it in host mode using the `ipv6 address autoconfig` command.
- IPv6 forwarding is permanently disabled on the management Ethernet interface so that it remains in Host mode and does not operate as a router regardless of the `ipv6 address autoconfig` setting.

If necessary, you can manually disable IPv6 processing on an interface so that the configured IPv6 addresses do not take effect. The IPv6 addresses take effect again when you re-enable IPv6.

If you disable IPv6 and configure a Layer (L2) interface in Layer (L3) mode, IPv6 is not automatically re-enabled on the interface. You must manually re-enable it.

A link-local address automatically generates when you re-enable IPv6 on an interface with the `ipv6 enable` command.

Disable and enable IPv6

```
OS10(config)# interface ethernet 1/1/8
OS10(conf-if-eth1/1/8)# ipv6 address 2111:dddd:0eee::22/64
OS10(conf-if-eth1/1/8)# no ipv6 address autoconfig
OS10(conf-if-eth1/1/8)# no ipv6 enable
OS10(conf-if-eth1/1/8)# ipv6 enable
```

Display IPv6 status

```
OS10# show interface ethernet 1/1/20
Ethernet 1/1/20 is up, line protocol is up
Hardware is Dell EMC Eth, address is ec:f4:bb:fb:fa:30
    Current address is ec:f4:bb:fb:fa:30
Pluggable media present, QSFP+ type is QSFP+ 40GBASE CR 1.0M
    Wavelength is 850
    Receive power reading is 0.0
Interface index is 17305562
Internet address is 20.20.20.1/24
Mode of IPv4 Address Assignment: MANUAL
Interface IPv6 oper status: Enabled
Link local IPv6 address: fe80::eef4:bbff:fefb:fa30/64
Global IPv6 address: 2020::1/64
...
```

```
OS10# show ipv6 interface brief
Interface Name    admin/protocol  IPV6 Address/Link-Local Address  IPv6 Oper Status
================================================================================
Ethernet 1/1/1:1  up   / up       fe80::eef4:bbff:fefb:f9f0/64
                                  2017::1/64                      Enabled
Ethernet 1/1/20   up   / up       fe80::eef4:bbff:fefb:fa30/64
                                  2020::1/64                      Enabled
Management 1/1/1  up   / up       fe80::eef4:bbff:fefb:f9ef/64    Enabled
Vlan 1            up   / up       fe80::eef4:bbff:fefb:fa59/64    Enabled
```

IPv6 addresses

An IPv6 address consists of a 48-bit global routing prefix, optional 16-bit subnet ID, and a 64-bit interface identifier in the extended universal identifier (EUI)-64 format.

IPv6 128-bit addresses are represented as a series of eight 16-bit hexadecimal fields separated by colons: x:x:x:x:x:x:x:x.

```
2001:0db8:0000:0000:0000:0000:1428:57a
```

Leading zeros in each field are optional. You can also use two colons (::) to represent successive hexadecimal fields of zeros, but you can use this short version only one time in each address:

```
2001:db8::1428:57ab
```

In the following example, all the addresses are valid and equivalent:

- `2001:0db8:0000:0000:0000:0000:1428:57ab`
- `2001:0db8:0000:0000:0000::1428:57ab`
- `2001:0db8:0:0:0:0:1428:57ab`
- `2001:0db8:0:0::1428:57ab`
- `2001:0db8::1428:57ab`
- `2001:db8::1428:57ab`

Write IPv6 networks using CIDR notation. An IPv6 network or subnet is a contiguous group of IPv6 addresses which must be a power of two. The initial bits of addresses, which are identical for all hosts in the network, are the network's prefix.

A network is denoted by the first address in the network and the size in bits of the prefix in decimal, separated with a slash. Because a single host is seen as a network with a 128-bit prefix, host addresses may be written with a following /128.

For example, `2001:0db8:1234::/48` stands for the network with addresses `2001:0db8:1234:0000:0000:0000:0000:0000` through `2001:0db8:1234:ffff:ffff:ffff:ffff:ffff`.

As soon as you assign an IPv6 address, IPv6 packet processing is enabled on an interface. You can manually disable and re-enable IPv6 processing on an interface configured with an IPv6 address using the `no ipv6 enable` and `ipv6 enable` commands.

To remove all IPv6 addresses from an interface, use the `no ipv6 address` command. To remove a specific IPv6 address, use the `ipv6 address ipv6-address/mask` command.

Link-local addresses

When an OS10 switch boots up, an IPv6 unicast link-local address automatically assigns to an interface using stateless configuration. A link-local address allows IPv6 devices on a local link to communicate without requiring a globally unique address. IPv6 reserves the address block FE80::/10 for link-local unicast addressing.

Global addresses

To enable stateless autoconfiguration of an IPv6 global address and set the interface to Host mode, use the `ipv6 address autoconfig` command. The router receives network prefixes in IPv6 router advertisements (RAs). An interface ID appends to the prefix. In Host mode, IPv6 forwarding is disabled.

The `no ipv6 address autoconfig` command disables IPv6 global address autoconfiguration, and sets the interface to Router mode with IPv6 forwarding enabled.

DHCP-assigned addresses

As an alternative to stateless autoconfiguration, you can enable a network host to obtain IPv6 addresses using a DHCP server via stateful autoconfiguration using the `ipv6 address dhcp` command. A DHCPv6 server uses a prefix pool to configure a network address on an interface. The interface ID automatically generates.

Manally configured addresses

An interface can have multiple IPv6 addresses. To configure an IPv6 address in addition to the link-local address, use the `ipv6 address ipv6-address/mask` command. Enter the full 128-bit IPv6 address, including the network prefix and a 64-bit interface ID.

You can also manually configure an IPv6 address by assigning:

- A network prefix with the EUI-64 parameter using the `ipv6 address ipv6-prefix eui64` command. A 64-bit interface ID automatically generates based on the MAC address.
- A link-local address to use instead of the link-local address that automatically configures when you enable IPv6 using the `ipv6 address link-local` command.

Configure IPv6 address

```
OS10(config)# interface ethernet 1/1/8
OS10(conf-if-eth1/1/8)# ipv6 address 2001:dddd:0eee::4/64
```

Configure network prefix

```
OS10(config)# interface ethernet 1/1/8
OS10(conf-if-eth1/1/8)# ipv6 address 2001:FF21:1:1::/64 eui64
```

Configure link-local address

```
OS10(config)# interface ethernet 1/1/8
OS10(conf-if-eth1/1/8)# ipv6 address FE80::1/64 link-local
```

Stateless autoconfiguration

When an interface comes up, OS10 uses stateless autoconfiguration to generate a unique link-local IPv6 address with a FE80::/64 prefix and an interface ID generated from the MAC address. To use stateless autoconfiguration to assign a globally unique address using a prefix received in router advertisements, use the `ipv6 address autoconfig` command.

Stateless autoconfiguration sets an interface in Host mode, and allows the interface connected to an IPv6 network to autoconfigure IPv6 addresses and communicate with other IPv6 devices on local links. A DHCP server is not required for automatic IPv6 interface

configuration. IPv6 devices on a local link send router advertisement (RA) messages in response to solicitation messages received at startup.

Perform stateless autoconfiguration of IPv6 addresses using:

Prefix advertisement	Routers use router advertisement messages to advertise the network prefix. Hosts append their interface-identifier MAC address to generate a valid IPv6 address.
Duplicate address detection	An IPv6 host node checks whether that address is used anywhere on the network using this mechanism before configuring its IPv6 address.
Prefix renumbering	Transparent renumbering of hosts in the network when an organization changes its service provider.

IPv6 provides the flexibility to add prefixes on RAs in response to a router solicitation (RS). By default, RA response messages are sent when an RS message is received. The system manipulation of IPv6 stateless autoconfiguration supports the router side only. Neighbor Discovery (ND) messages advertise so the neighbor can use the information to auto-configure its address. Received ND messages are not used to create an IPv6 address.

Inconsistencies in RA values between routers are logged. The values checked for consistency include:

- Current hop limit
- M and O flags
- Reachable time
- Retransmission timer
- MTU options
- Preferred and valid lifetime values for the same prefix

The router redirect functionality in the NDP is similar to IPv4 router redirect messages. NDP uses ICMPv6 redirect messages (Type 137) to inform nodes that a better router exists on the link.

Neighbor Discovery

The IPv6 NDP determines if neighboring IPv6 devices are reachable and receives the IPv6 addresses of IPv6 devices on local links. Using the link-layer and global prefixes of neighbor addresses, OS10 performs stateless autoconfiguration of IPv6 addresses on interfaces.

ICMPv6 RA messages advertise the IPv6 addresses of IPv6-enabled interfaces and allow a router to learn of any address changes in IPv6 neighbors. By default, RAs are disabled on an interface.

Prerequisites

To enable RA messages, the switch must be in Router mode with IPv6 forwarding enabled and stateless autoconfiguration disabled using the `no ipv6 address autoconfig` command.

Enable router advertisement messages

1 Enable IPv6 neighbor discovery and sending ICMPv6 RA messages in Interface mode.

```
ipv6 nd send-ra
```

2 (Optional) Configure IPv6 neighbor discovery options in Interface mode.

- `ipv6 nd hop-limit` *hops* — (Optional) Sets the hop limit advertised in RA messages and included in IPv6 data packets sent by the router, from 0 to 255; default 64. `0` indicates that no hop limit is specified by the router.
- `ipv6 nd managed-config-flag` — (Optional) Sent in RA messages to tell hosts to use stateful address autoconfiguration, such as DHCPv6, to obtain IPv6 addresses.
- `ipv6 nd max-ra-interval` *seconds* — (Optional) Sets the maximum time interval for sending RA messages, from 4 to 1800 seconds; default 600.
- `ipv6 nd mtu` *number* — (Optional) Sets the maximum transmission unit (MTU) used in RA messages on the link, from 1280 to 65535 bytes; default 1500. By default, no MTU setting is included in RA messages.

- `ipv6 nd other-config-flag` — (Optional) Tells hosts to use stateful autoconfiguration to obtain nonaddress-related information.
- `ipv6 nd ra-lifetime` *`seconds`* — (Optional) Sets the lifetime of a default router in RA messages, from 0 to 9000 milliseconds; default 3 times the `max-ra-interval` setting. `0` indicates that this router is not used as a default router.
- `ipv6 nd reachable-time` *`milliseconds`* — (Optional) Sets the advertised time the router sees that a neighbor is up after it receives neighbor reachability confirmation, from 0 to 3600000 milliseconds; default 0. `0` indicates that no reachable time is sent in RA messages.
- `ipv6 nd retrans-timer` *`seconds`* — (Optional) Sets the time between retransmitting neighbor solicitation messages, from 100 to 4292967295 milliseconds. By default, no retransmit timer is configured.

3 Configure the IPv6 prefixes that are advertised by IPv6 neighbor discovery in Interface mode.

```
ipv6 nd prefix {ipv6-prefix | default} [no-advertise] [no-autoconfig] [no-rtr-address]
[off-link] [lifetime {valid-lifetime seconds | infinite}
{preferred-lifetime seconds | infinite}]
```

- `ipv6-prefix` — Enter an IPv6 prefix in *x:x::y/mask* format to include the prefix in RA mesages. Include prefixes that are not already in the subnets configured on the interface.
- `default` — Configure the prefix parameters advertised in all subnets configured on the interface.
- `no-advertise` — (Optional) Do not advertise the specified prefix. By default, all prefixes in configured subnets are advertised.
- `no-autoconfig` — (Optional) Sets `AdvAutonomous` to `Off` for the specified prefix in the radvd.conf file. This setting tells hosts to not use this prefix for address autoconfiguration. By default, `AdvAutonomous` is `On`.
- `no-rtr-address` — (Optional) Sets `AdvRouterAddr` to `Off` for the prefix in the radvd.conf file. The `Off` setting tells hosts to not use the advertising router address for on-link determination. By default, `AdvRouterAddr` is `On`.
- `off-link` — (Optional) Sets `AdvOnLink` to `Off` for the prefix in the radvd.conf file. The `Off` setting tells hosts to not use this prefix for on-link determination. By default, `AdvOnLink` is `On`.
- `lifetime {valid-lifetime` *`seconds`* `| infinite}` — (Optional) Sets `AdvValidLifetime` in seconds for the prefix in the radvd.conf file. The prefix is valid for on-link determination only for the specified lifetime. The default is 86400 seconds (1 day). The `infinite` setting allows the prefix to be valid for on-link determination with no time limit.
- `lifetime {preferred-lifetime` *`seconds`* `| infinite}` — (Optional) Sets `AdvPreferredLifetime` in seconds for the prefix in the radvd.conf file. IPv6 addresses generated from the prefix using stateless autoconfiguration remain preferred for the configured lifetime. The default is 14400 seconds (4 hours). The `infinite` setting allows addresses that are autoconfigured using the prefix to be preferred with no time limit.

By default, all prefixes configured in IPv6 addresses on an interface are advertised. To modify the default values advertised for interface subnet prefixes, use the `ipv6 nd prefix default` command and specify new default settings.

On-link determination is the process used to forward IPv6 packets to a destination IPv6 address.

Configure neighbor discovery

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ipv6 nd mtu 1500
OS10(conf-if-eth1/1/1)# ipv6 nd send-ra
```

Configure advertised IPv6 prefixes

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ipv6 nd prefix default lifetime infinite infinite
OS10(conf-if-eth1/1/1)# ipv6 nd prefix 2002::/64
```

Duplicate address discovery

To determine if an IPv6 unicast address is unique before assigning it to an interface, an OS10 switch sends a neighbor solicitation message. If the process of duplicate address discovery (DAD) detects a duplicate address in the network, the address does not configure on the interface. DAD is enabled by default.

By default, IPv6 is not disabled when a duplicate address is detected. Only the duplicate address is not applied. Other IPv6 addresses are still active on the interface.

To disable IPv6 on an interface when a duplicate link-local address is detected, use the `ipv6 nd dad disable-ipv6-on-failure` command. To re-enable IPv6 after you resolve a duplicate link-local address, enter `no ipv6 enable`, then the `ipv6 enable` command.

- Disable or re-enable IPv6 duplicate address discovery in Interface mode.
  ```
  ipv6 nd dad {disable | enable}
  ```
- Disable IPv6 on an interface if a duplicate link-local address is discovered in Interface mode.
  ```
  ipv6 nd dad disable-ipv6-on-dad-failure
  ```

Disable duplicate address discovery

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ipv6 nd dad disable
```

Disable IPv6 for duplicate link-local address

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ipv6 nd dad disable-ipv6-on-dad-failure
```

Static IPv6 routing

To define an explicit route between two IPv6 networking devices, configure a static route on an interface. Static routing is useful for smaller networks with only one path to an outside network, or to provide security for certain traffic types in a larger network.

- Enter the static routing information including the IPv6 address and mask in x:x:x:x::x format in CONFIGURATION mode. The length is from 0 to 64.
  ```
  ipv6 route ipv6-prefix/mask {next-hop | interface interface [route-preference]}
  ```
  - *next-hop* — Enter the next-hop IPv6 address in x:x:x:x::x format.
  - `interface` *interface* — Enter the interface type then the slot/port or number information.
  - *route-preference* — (Optional) Enter a route-preference range, from 1 to 255.

After you configure a static IPv6 route, configure the forwarding router's address on the interface. The IPv6 neighbor interface must have an IPv6 address configured.

Configure IPv6 static routing and view configuration

```
OS10(config)# ipv6 route 2111:dddd:0eee::22/128 2001:db86:0fff::2
OS10(config)# do show ipv6 route static
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF,IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, > - non-active route
Gateway of last resort is not set
  Destination            Gateway                 Dist/Metric   Last Change
------------------------------------------------------------------------------
  S 2111:dddd:eee::22/12via 2001:db86:fff::2 ethernet1/1/1  1/1  00:01:24
```

IPv6 destination unreachable

By default, when no matching entry for an IPv6 route is found in the IPv6 routing table, a packet drops and no error message is sent. You can enable the capability to send an `IPv6 destination unreachable` error message to the source without dropping the packet.

Enable IPv6 unreachable destination messaging

```
OS10(config)# interface ethernet 1/1/8
OS10(conf-if-eth1/1/8)# ipv6 unreachables
```

IPv6 hop-by-hop options

A hop-by-hop header extension in an IPv6 packet contains options that are processed by all IPv6 routers in the packet's path. By default, hop-by-hop header options in an IPv6 packet do not process locally. To enable local processing of IPv6 hop-by-hop options on an interface, use the `ipv6 hop-by-hop` command.

Enable IPv6 hop-by-hop options forwarding

```
OS10(config)# interface ethernet 1/1/8
OS10(conf-if-eth1/1/8)# ipv6 hop-by-hop
```

View IPv6 information

To view IPv6 configuration information, use the `show ipv6 route` command. To view IPv6 address information, use the `show address ipv6` command.

View IPv6 connected information

```
OS10# show ipv6 route connected
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF,IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, > - non-active route
Gateway of last resort is not set
Destination       Gateway                 Dist/Metric  Last Change
----------------------------------------------------------------
C 2001:db86::/32  via 2001:db86:fff::1 ethernet1/1/1  0/0   00:03:24
```

View IPv6 static information

```
OS10# show ipv6 route static
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF,IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, > - non-active route
Gateway of last resort is not set
Destination         Gateway                 Dist/Metric  Last Change
----------------------------------------------------------------
S  2111:dddd:eee::22/12via 2001:db86:fff::2 ethernet1/1/1  1/1 00:01:24
```

IPv6 commands

clear ipv6 neighbors

Deletes all entries in the IPv6 neighbor discovery cache or neighbors of a specific interface. Static entries are not removed.

Syntax

```
clear ipv6 neighbors [vrf vrf-name] [ipv6-address | interface | virtual-network vn-id | all]
```

Parameters

- `vrf vrf-name` — (Optional) Enter `vrf` then the name of the VRF to clear the neighbor corresponding to that VRF. If you do not specify this option, the neighbors in the default VRF clear.
- `ipv6-address` — Enter the IPv6 address of the neighbor in the x:x:x:x::x format to remove a specific IPv6 neighbor. The :: notation specifies successive hexadecimal fields of zero.
- `interface interface` — To remove all neighbor entries learned on a specific interface, enter the keyword interface then the interface type and slot/port or number information of the interface:
  - For a 10-Gigabit Ethernet interface, enter `TenGigabitEthernet` then the slot/port/subport[/subport] information.
  - For a 40-Gigabit Ethernet interface, enter `fortyGigE` then the slot/port information.
  - For a port channel interface, enter `port-channel` then a number.
  - For a VLAN interface, enter `vlan` then a number from 1 to 4093.
  - `virtual-network vn-id` — For a virtual network, enter `virtual-network` then the ID of the network.

Defaults None.

Command Mode EXEC

Usage Information The `no` version of this command resets the value to the default.

Example

Supported Releases 10.4.1.0 or later or later

clear ipv6 route

Clears routes from the IPv6 routing table.

Syntax

```
clear ipv6 route [vrf vrf-name] {* | A::B/mask}
```

Parameters

- `vrf vrf-name` — (Optional) Enter `vrf` then the name of the VRF to clear the IPv6 routes corresponding to that VRF.
- `*` — Clears all routes and refreshes the IPv6 routing table. Traffic flow for all the routes in the switch is affected.
- `A::B/mask` — Removes the IPv6 route and refreshes the IPv6 routing table. Traffic flow in the switch is affected only for the specified route.

Default Not configured

Command Mode EXEC

Usage Information This command does not remove the static routes from the routing table.

Example

```
OS10# clear ipv6  route *
```

Supported Releases 10.3.0E or later

ipv6 address

Configures a global unicast IPv6 address on an interface.

Syntax `ipv6 address ipv6-address/prefix-length`

Parameters *ipv6-address/prefix-length* — Enter a full 128-bit IPv6 address with the network prefix length, including the 64-bit interface identifier.

Defaults None

Command Mode INTERFACE

Usage Information An interface can have multiple IPv6 addresses. To configure an IPv6 address in addition to the link-local address, use the `ipv6 address ipv6-address/mask` command and specify the complete 128-bit IPv6 address. To configure a globally unique IPv6 address by entering only the network prefix and length, use the `ipv6 address ipv6-prefix/prefix-length eui-64` command.

The `no` version of this command removes the IPv6 address on the interface.

Example

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ipv6 address 2111:dddd:0eee::22/64
```

Supported Releases 10.3.0E or later

ipv6 address autoconfig

Acquires global IPv6 addresses by using the network prefix obtained from RAs.

Syntax `ipv6 address autoconfig`

Parameters None

Defaults Disabled except on the management interface

Command Mode INTERFACE

Usage Information

- This command sets an interface in Host mode to perform IPv6 stateless auto-configuration by discovering prefixes on local links, and adding an EUI-64 based interface identifier to generate each IPv6 address. The command disables IPv6 forwarding. Addresses are configured depending on the prefixes received in RA messages.
- The `no` version of this command disables IPv6 address autoconfiguration, resets the interface in Router mode, and re-enables IPv6 forwarding.

Example

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# ipv6 address autoconfig
OS10(conf-if-eth1/1/1)#
```

Supported Releases 10.3.0E or later

ipv6 address dhcp

Enables DHCP client operations on the interface.

Syntax	`ipv6 address dhcp`
Parameters	None
Defaults	None
Command Mode	INTERFACE
Usage Information	The `no` version of this command disables DHCP operations on the interface.
Example	`OS10(config)# interface mgmt 1/1/1` `OS10(conf-if-ma-1/1/1)# ipv6 address dhcp`
Supported Releases	10.3.0E or later

ipv6 enable

Enables and disables IPv6 forwarding on an interface configured with an IPv6 address.

Syntax	`ipv6 enable`
Parameters	None
Defaults	None
Command Mode	INTERFACE
Usage Information	Use this command to disable and re-enable IPv6 forwarding on an interface for security purposes or to recover from a duplicate address discovery (DAD) failure. The `no` version of this command disables IPv6 forwarding.

Example	`OS10(config)# interface ethernet 1/1/1`
	`OS10(conf-if-eth1/1/1)# ipv6 address 2111:dddd:0eee::22/128`
	`OS10(conf-if-eth1/1/1)# no ipv6 enable`
	`OS10(conf-if-eth1/1/1)# ipv6 enable`

Supported Releases	10.3.0E or later

ipv6 address eui-64

Configures a global IPv6 address on an interface by entering only the network prefix and length.

Syntax	`ipv6 address` *`ipv6-prefix/prefix-length`* `eui-64`
Parameters	*ipv6-prefix* — Enter an IPv6 prefix in *x:x::y/mask* format.
Defaults	None
Command Mode	INTERFACE
Usage Information	Use this command to manually configure an IPv6 address in addition to the link-local address generated with stateless autoconfiguration. Specify only the network prefix and length. The 64-bit interface ID automatically computes from the MAC address. This command enables IPv6 processing on the interface. The `no` version of this command removes the IPv6 address configuration.

Example	`OS10(config)# interface mgmt 1/1/1` `OS10(conf-if-ma-1/1/1)# ipv6 address 2111:dddd:0eee::/64 eui-64`
Supported Releases	10.4.0E(R1) or later

ipv6 address link-local

Configures a link-local IPv6 address on the interface to use instead of the link-local address that is automatically configured with stateless autoconfiguration.

Syntax	`ipv6 address` *`ipv6-prefix`* `link-local`
Parameters	*`ipv6-prefix`* — Enter an IPv6 prefix in *x:x::y/mask* format.
Defaults	None
Command Mode	INTERFACE
Usage Information	• An interface can have only one link-local address. By default, an IPv6 link-local address automatically generates with a MAC-based EUI-64 interface ID when a router boots up and IPv6 is enabled. Use this command to manually configure a link-local address to replace the autoconfigured address. For example, to configure a more user-friendly link-local address, replace `fe80::eef4:bbff:fefb:fa30/64` with `fe80::1/64`. • The `no` version of this command removes the specified link-local address.
Example	`OS10(config)# interface mgmt 1/1/1` `OS10(conf-if-ma-1/1/1)# ipv6 address 2111:dddd:0eee::22/64 link-local`
Supported Releases	10.4.0E(R1) or later

ipv6 hop-by-hop

Enables and disables processing hop-by-hop options in IPv6 packet headers.

Syntax	`ipv6 hop-by-hop`
Parameters	None
Defaults	Hop-by-hop header options in an IPv6 packet do not process on an interface.
Command Mode	INTERFACE
Usage Information	• Use this command to enable local processing of IPv6 packets with hop-by-hop options in conformance with the RFC 8200, IPv6 Specification. • The `no` version of this command disables IPv6 processing of hop-by-hop header options.
Example: Disable hop-by-hop option processing	`OS10(config)# interface ethernet 1/2/3` `OS10(conf-if-eth1/2/3)# no ipv6 hop-by-hop`
Supported Releases	10.4.0E(R1) or later

ipv6 nd dad

Disables or re-enables IPv6 duplicate address discovery (DAD).

Syntax

```
ipv6 nd dad {disable | enable | disable-ipv6-on-dad-failure}
```

Parameters

- `disable` — Disable duplicate address discovery on the interface.
- `enable` — Re-enable IPv6 duplicate address discovery if you have disabled it.
- `disable-ipv6-on-dad-failure` — Enable duplicate address discovery on the existing autoconfigured link-local address.

Defaults

Duplicate address discovery is enabled on an interface.

Command Mode

INTERFACE

Usage Information

- An OS10 switch sends a neighbor solicitation message to determine if an autoconfigured IPv6 unicast link-local address is unique before assigning it to an interface. If the process of duplicate address discovery (DAD) detects a duplicate address in the network, the link-local address does not configure. Other IPv6 addresses are still active on the interface.
- By default, DAD does not disable IPv6 if a duplicate link-local address is detected in the network. To disable IPv6 on an interface when a duplicate link-local address is detected, use the `ipv6 nd dad disable-ipv6-on-failure` command.

Example: Disable DAD

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ipv6 nd dad disable
```

Example: Enable DAD on link-local address

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ipv6 nd dad disable-ipv6-on-dad-failure
```

Supported Releases

10.4.0E(R1) or later

ipv6 nd hop-limit

Sets the hop limit advertised in RA messages and included in IPv6 data packets sent by the router.

Syntax

```
ipv6 nd hop-limit hops
```

Parameters

- `hop-limit hops` — Enter the maximum number of hops allowed for RA messages, from 0 to 255.

Defaults

64 hops

Command Mode

INTERFACE

Usage Information

The configured hop limit is advertised in RA messages and included in IPv6 data packets sent by the router. `0` indicates that no hop limit is specified by the router.

Example

```
OS10(config)# interface ethernet 1/2/3
OS10(conf-if-eth1/2/3)# ipv6 nd hop-limit 100
```

Supported Releases

10.4.0E(R1) or later

ipv6 nd managed-config-flag

Sends RA messages that tell hosts to use stateful address autoconfiguration, such as DHCPv6, to obtain IPv6 addresses.

Syntax `ipv6 nd managed-config-flag`

Parameters None

Defaults Not configured

Command Mode INTERFACE

Usage Information The `no` version of this command disables the `managed-config-flag` option in RA messages.

Example

```
OS10(config)# interface ethernet 1/2/3
OS10(conf-if-eth1/2/3)# ipv6 nd managed-config-flag
```

Supported Releases 10.4.0E(R1) or later

ipv6 nd max-ra-interval

Sets the maximum time interval between sending RA messages.

Syntax `ipv6 nd max-ra-interval` *seconds*

Parameters

- `max-ra-interval` *seconds* — Enter a time interval in seconds, from 4 to 1800.

Defaults 600 seconds

Command Mode INTERFACE

Usage Information The `no` version of this command restores the default time interval used to send RA messages.

Example

```
OS10(config)# interface ethernet 1/2/3
OS10(conf-if-eth1/2/3)# ipv6 nd max-ra-interval 300
```

Supported Releases 10.4.0E(R1) or later

ipv6 nd mtu

Sets the maximum transmission unit (MTU) used on a local link in RA messages.

Syntax `ipv6 nd mtu` *number*

Parameters

- `mtu` *number* — Enter the MTU size in bytes, from 1280 to 65535.

Defaults 1500 bytes

Command Mode INTERFACE

Usage Information The `no` version of this command restores the default MTU value advertised in RA messages.

Example

```
OS10(config)# interface ethernet 1/2/3
OS10(conf-if-eth1/2/3)# ipv6 nd mtu 2500
```

Supported Releases 10.4.0E(R1) or later

ipv6 nd other-config-flag

Sends RA messages that tell hosts to use stateful autoconfiguration to obtain nonaddress-related information.

Syntax `ipv6 nd other-config-flag`

Parameters None

Defaults Not configured

Command Mode INTERFACE

Usage Information The `no` version of this command disables the `other-config-flag` option in RA messages.

Example

```
OS10(config)# interface ethernet 1/2/3
OS10(conf-if-eth1/2/3)# ipv6 nd other-config-flag
```

Supported Releases 10.4.0E(R1) or later

ipv6 nd prefix

Configures the IPv6 prefixes that are included in messages to neighboring IPv6 routers.

Syntax `ipv6 nd prefix {`*`ipv6-prefix`*` | default} [no-advertise] [no autoconfig] [no-rtr-address] [off-link] [lifetime {valid-lifetime `*`seconds`*` | infinite} {preferred-lifetime `*`seconds`*` | infinite}]`

Parameters

- `ipv6-prefix` — Enter an IPv6 prefix in *x:x::y/mask* format to include the prefix in RA mesages. Include prefixes that are not already in the subnets on the interface.
- `default` — Configure the prefix parameters advertised in all subnets configured on the interface.
- `no-advertise` — (Optional) Do not advertise the specified prefix. By default, all prefixes in configured subnets advertise.
- `no-autoconfig` — (Optional) Sets `AdvAutonomous` to `Off` for the specified prefix in the radvd.conf file. This setting tells hosts to not use this prefix for address autoconfiguration. By default, `AdvAutonomous` is `On`.
- `no-rtr-address` — (Optional) Sets `AdvRouterAddr` to `Off` for the prefix in the radvd.conf file. The `Off` setting tells hosts to not use the advertising router's address for on-link determination. By default, `AdvRouterAddr` is `On`.
- `off-link` — (Optional) Sets `AdvOnLink` to `Off` for the prefix in the radvd.conf file. The `Off` setting tells hosts to not use this prefix for on-link determination. By default, `AdvOnLink` is `On`.
- `lifetime {valid-lifetime` *`seconds`* `| infinite}` — (Optional) Sets `AdvValidLifetime` in seconds for the prefix in the radvd.conf file. The prefix is valid for on-link determination only for the specified lifetime. The default is 86400 seconds (1 day). The `infinite` setting allows the prefix to be valid for on-link determination with no time limit.
- `lifetime {preferred-lifetime` *`seconds`* `| infinite}` — (Optional) Sets `AdvPreferredLifetime` in seconds for the prefix in the radvd.conf file. IPv6 addresses generated from the prefix using stateless autoconfiguration remain preferred for the configured lifetime. The default is 14400 seconds (4 hours). The `infinite` setting allows addresses that are autoconfigured using the prefix to be preferred with no time limit.

Defaults All prefixes in IPv6 subnets configured on an interface advertise.

Command Mode INTERFACE

Usage Information

- By default, all prefixes configured in IPv6 addresses on an interface advertise. To advertise all default parameters in the subnet prefixes on an interface, enter the `default` keyword.

- If you configure a prefix with valid or preferred lifetime values, the `ipv6 nd prefix default no autoconfig` command does not apply the default prefix values.
- On-link determination is used to forward IPv6 packets to a destination IPv6 address.

Examples

Enable router advertisements

```
OS10(conf-if-eth1/1/1)# ipv6 address 2001:0db8:2000::1/64
OS10(conf-if-eth1/1/1)# ipv6 nd send-ra
```

Change default settings for interface subnet prefixes

```
OS10(conf-if-eth1/1/1)# ipv6 nd prefix default lifetime infinite infinite
```

Disable advertising an interface subnet prefix

```
OS10(conf-if-eth1/1/1)# ipv6 nd prefix 2001:0db8:2000::/64 no-advertise
```

Advertise prefix for which there is no interface address

```
OS10(conf-if-eth1/1/1)# ipv6 nd prefix 2001:0db8:3000::/64 no-autoconfig
```

Supported Releases 10.4.0E(R1) or later

ipv6 nd ra-lifetime

Sets the lifetime of the default router in RA messages.

Syntax `ipv6 nd ra-lifetime seconds`

Parameters

- `ra-lifetime seconds` — Enter a lifetime value in milliseconds, from 0 to 9000 milliseconds.

Defaults Three times the `max-ra-interval` value

Command Mode INTERFACE

Usage Information The `no` version of this command restores the default lifetime value. `0` indicates that this router is not used as the default router.

Example

```
OS10(config)# interface ethernet 1/2/3
OS10(conf-if-eth1/2/3)# ipv6 nd max-ra-interval 300
```

Supported Releases 10.4.0E(R1) or later

ipv6 nd reachable-time

Sets the advertised time the router sees a neighbor to be up after it receives a reachability confirmation.

Syntax `ipv6 nd reachable-time milliseconds`

Parameters

- `reachable-time milliseconds` — Enter the reachable time in milliseconds, from 0 to 3600000.

Defaults 0

Command Mode INTERFACE

Usage Information The `no` version of this command restores the default reachable time. `0` indicates that no reachable time is sent in RA messages.

Example

```
OS10(config)# interface ethernet 1/2/3
OS10(conf-if-eth1/2/3)# ipv6 nd reachable-time 1000
```

Supported Releases 10.4.0E(R1) or later

ipv6 nd retrans-timer

Sets the time between retransmitting neighbor solicitation messages.

Syntax `ipv6 nd retrans-timer seconds`

Parameters

- `retrans-timer seconds` — Enter the retransmission time interval in milliseconds, from 100 to 4292967295.

Defaults Not configured

Command Mode INTERFACE

Usage Information The `no` version of this command disables the configured retransmission timer.

Example

```
OS10(config)# interface ethernet 1/2/3
OS10(conf-if-eth1/2/3)# ipv6 nd retrans-timer 1000
```

Supported Releases 10.4.0E(R1) or later

ipv6 nd send-ra

Enables sending ICMPv6 RA messages.

Syntax `ipv6 nd send-ra`

Parameters None

Defaults RA messages are disabled.

Command Mode INTERFACE

Usage Information

- Using ICMPv6 RA messages, the Neighbor Discovery Protocol (NDP) advertises the IPv6 addresses of IPv6-enabled interfaces and learns of any address changes in IPv6 neighbors. Before you enable sending RA messages, the switch must be in Router mode with IPv6 forwarding enabled and stateless autoconfiguration disabled `no ipv6 address autoconfig` command.
- The `no` version command disables RA messages.

Example

```
OS10(config)# interface ethernet 1/2/3
OS10(conf-if-eth1/2/3)# ipv6 nd send-ra
```

Supported Releases 10.4.0E(R1) or later

ipv6 route

Configures a static IPv6 static route.

Syntax `ipv6 route [dst-vrf vrf-name] ipv6-prefix mask {next-hop | interface interface-type [route-preference]}`

Parameters

- `dst-vrf` *`vrf-name`* — (Optional) Enter `vrf` then the name of the VRF to install IPv6 routes in that VRF.
- *`ipv6-prefix`* — Enter the IPv6 address in x:x:x:x::x format.
- *`mask`* — Enter the mask in slash prefix-length /x format.
- *`next-hop`* — Enter the next-hop IPv6 address in x:x:x:x::x format.
- `interface` *`interface-type`* — Enter the interface type then the slot/port or number information. The interface types supported are: Ethernet, port-channel, VLAN, and Null.
- *`route-preference`* — (Optional) Enter a route-preference range, from 1 to 255.

Default Not configured

Command Mode CONFIGURATION

Usage Information

- When the interface fails, the system withdraws the route. The route reinstalls when the interface comes back up. When a recursive resolution breaks, the system withdraws the route. The route reinstalls when the recursive resolution is satisfied. After you create an IPv6 static route interface, if you do not assign an IP address to a peer interface, you must manually ping the peer to resolve the neighbor information.
- The `no` version of this command deletes the IPv6 route configuration.

Example

```
OS10(config)# ipv6 route 2111:dddd:0eee::22/128 2001:db86:0fff::2
OS10(config)# ipv6 route 2111:dddd:0eee::22/128 interface null 0
```

Supported Releases 10.2.0E or later

ipv6 unreachables

Enables generating error messages on an interface for IPv6 packets with unreachable destinations.

Syntax `ipv6 unreachables`

Parameters None

Defaults ICMPv6 unreachable messages are not sent.

Command Mode INTERFACE

Usage Information

- By default, when no matching entry for an IPv6 route is found in the IPv6 routing table, the packet drops and no error message is sent. Use this command to enable sending an `IPv6 destination unreachable` error message to the source without dropping the packet.
- The `no` version of this command disables generating unreachable destination messages.

Example

```
OS10(config)# interface ethernet 1/2/3
OS10(conf-if-eth1/2/3)# ipv6 unreachables
```

Supported Releases 10.4.0E(R1) or later

show ipv6 neighbors

Displays IPv6 discovery information. Entering the command without options shows all IPv6 neighbor addresses stored on the control processor (CP).

Syntax `show ipv6 neighbors [vrf` *`vrf-name`*`] [ipv6-address| interface interface]`

Parameters

- `vrf vrf-name` — (Optional) Enter `vrf` then the name of the VRF to display the neighbors corresponding to that VRF. If you do not specify this option, neighbors corresponding to the default VRF display.
- `ipv6-address` — Enter the IPv6 address of the neighbor in the x:x:x:x::x format. The :: notation specifies successive hexadecimal fields of zero.
- `interface interface` — Enter `interface` then the interface type and slot/port or number information:
  - For a 10-Gigabit Ethernet interface, enter `TenGigabitEthernet` then the slot/port/subport[/subport] information.
  - For a 40-Gigabit Ethernet interface, enter `fortyGigE` then the slot/port information.
  - For a port channel interface, enter `port-channel` then a number.
  - For a VLAN interface, enter `vlan` then a number from 1 to 4093.

Defaults None.

Command Mode EXEC

Usage Information The `no` version of this command resets the value to the default.

Example

```
OS10# show ipv6 neighbors
IPv6 Address     Hardware Address    State   Interface       VLAN
-------------------------------------------------------------------
1001:db8:a1::2   00:c5:05:02:12:91   REACH   ethernet1/1/5   12
1001:db8:a1::f   00:f5:50:02:54:75   REACH   port-channel5   12
200::2           00:c5:05:02:12:91   STALE   ethernet1/1/10
400::f           00:f5:50:02:54:75   REACH   port-channel20
```

Supported Releases 10.4.1.0 or later or later

show ipv6 route

Displays IPv6 routes.

Syntax `show ipv6 route [vrf vrf-name] [all | bgp | connected | static | A::B/mask | summary]`

Parameters

- `vrf vrf-name` — (Optional) Enter `vrf` then the name of the VRF to display IPv6 routes corresponding to that VRF. If you do not specify this option, routes corresponding to the default VRF display.
- `all`—(Optional) Displays all routes including nonactive routes.
- `bgp`—(Optional) Displays BGP route information.
- `connected`—(Optional) Displays only the directly connected routes.
- `static`—(Optional) Displays all static routes.
- *A::B/mask*—(Optional) Enter the IPv6 destination address and mask.

  NOTE: This option works only for the exact prefix and the mask length.
- `summary`—(Optional) Displays the IPv6 route summary.

Default Not configured

Command Mode EXEC

Usage Information None

Example (All)

```
OS10# show ipv6 route all
Codes: C - connected
```

```
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF,IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, > - non-active route
Gateway of last resort is not set
  Destination     Gateway        Dist/Metric     Last Change
------------------------------------------------------------
```

Example (Connected)

```
OS10# show ipv6 route connected
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF,IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, > - non-active route
Gateway of last resort is not set
Destination         Gateway            Dist/Metric      Last Change
------------------------------------------------------------------
C    2001:db86::/32 via 2001:db86:fff::1 ethernet1/1/1  0/0 00:03:24
```

Example (Summary)

```
OS10# show ipv6 route summary
Route Source        Active Routes  Non-Active Routes
Ospf                0              0
Bgp                 0              0
Connected           0              0
Static              0              0
Ospf Inter-area     0              0
NSSA External-1     0              0
NSSA External-2     0              0
Ospf External-1     0              0
Ospf External-2     0              0
Bgp Internal        0              0
Bgp External        0              0
Ospf Intra-area     0              0
Total               0              0
```

Supported Releases 10.2.0E or later

show ipv6 interface brief

Displays IPv6 interface information.

Syntax `show ipv6 interface brief`

Parameters `brief` — Displays a brief summary of IPv6 interface information.

Defaults None

Command Mode EXEC

Usage Information Use the `do show ipv6 interface brief` command to view IPv6 interface information in other modes.

Example (Brief)

```
OS10# show ipv6 interface brief

Interface       admin/  IPV6 Address/             IPv6 Oper
Name          protocol  Link-Local Address           Status
============================================================
Management 1/1/1 up/up  fe80::20c:29ff:fe54:c852/64  Enabled
Vlan 1          up/up   fe80::20c:29ff:fe54:c8bc/64  Enabled
Ethernet 1/1/2  up/up   fe80::20c:29ff:fe54:c853/64
                        100::1/64
                   1001:1:1:1:20c:29ff:fe54:c853/64  Enabled
```

```
Ethernet 1/1/3   up/up    fe80::4/64
                           3000::1/64
                           4000::1/64                    Disabled
Ethernet 1/1/4   up/up    fe80::4/64
                           4::1/64
                           5::1/64                       Enabled
```

Supported Releases 10.2.0E or later or later

Open shortest path first

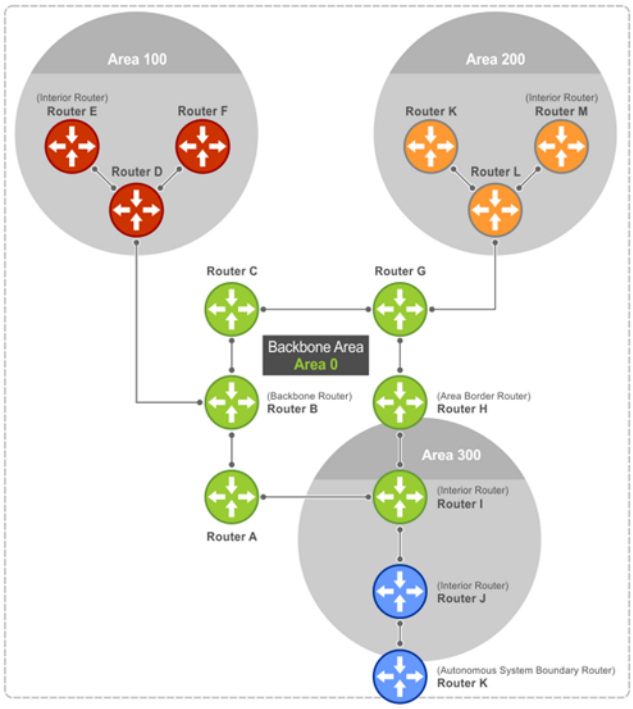
OSPF routing is a link-state routing protocol that allows sending link-state advertisements (LSAs) to all other routers within the same autonomous system (AS) area. OSPF LSAs include information about attached interfaces, metrics used, and other attributes. OSPF routers accumulate link-state information, and use the shortest path first (SPF) algorithm to calculate the shortest path to each node.

Autonomous system areas

OSPF operates in a hierarchy. The largest entity within the hierarchy is the autonomous system (AS). The AS is a collection of networks under a common administration that share a common routing strategy. OSPF is an intra-AS, Interior Gateway Routing Protocol (IGRP) that receives routes from and sends routes to other AS.

You can divide an AS into several areas, which are groups of contiguous networks and attached hosts administratively grouped. Routers with multiple interfaces can participate in multiple areas. These routers, called area border routers (ABRs), maintain separate databases for each area. Areas are a logical grouping of OSPF routers that an integer or dotted-decimal number identifies.

Areas allow you to further organize routers within the AS with one or more areas within the AS. Areas allow subnetworks to *hide* within the AS—minimizing the size of the routing tables on all routers. An area within the AS may not see the details of another area's topology. An area number or the router's IP address identifies AS areas.



Areas, networks, and neighbors

The backbone of the network is Area 0, also called Area 0.0.0.0, the core of any AS. All other areas must connect to Area 0. An OSPF backbone distributes routing information between areas. It consists of all area border routers and networks not wholly contained in any area and their attached routers.

The backbone is the only area with a default area number. You configure all other areas Area ID. If you configure two nonbackbone areas, you must enable the B bit in OSPF. Routers, A, B, C, G, H, and I are the backbone, see Autonomous system areas.

- A stub area (SA) does not receive external route information, except for the default route. These areas do receive information from interarea (IA) routes.
- A not-so-stubby area (NSSA) can import AS external route information and send it to the backbone as type-7 LSA.
- Totally stubby areas are also known as no summary areas.

Configure all routers within an assigned stub area as stubby and do not generate LSAs that do not apply. For example, a Type 5 LSA is intended for external areas and the stubby area routers may not generate external LSAs. A virtual link cannot traverse stubby areas.

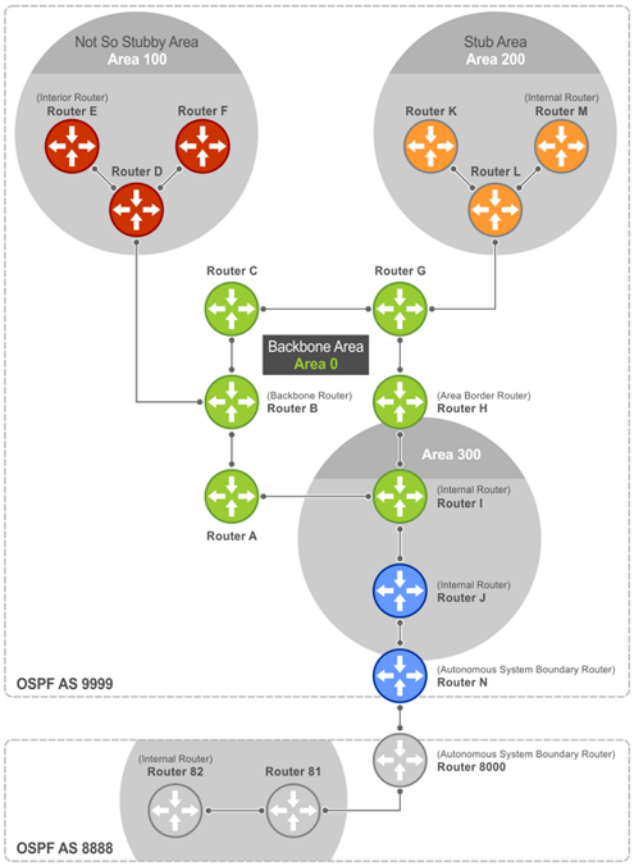
Networks and neighbors

As a link-state protocol, OSPF sends routing information to other OSPF routers concerning the state of the links between them. The Up or Down state of those links is important. Routers that share a link become neighbors on that segment. OSPF uses the `hello` protocol as a neighbor discovery and `keepalive` mechanism. After two routers are neighbors, they may proceed to exchange and synchronize their databases, which creates an adjacency.

Router types

Router types are attributes of the OSPF process—multiple OSPF processes may run on the same router. A router connected to more than one area, receiving routing from a BGP process connected to another AS, acts as both an area border router and an autonomous system border router.

Each router has a unique ID, written in decimal A.B.C.D format. You do not have to associate the router ID with a valid IP address. To make troubleshooting easier, ensure the router ID is identical to the router's IP address.



Backbone router	A backbone router (BR) is part of the OSPF Backbone, Area 0, and includes all ABRs. The BR includes routers connected only to the backbone and another ABR, but are only part of Area 0—shown as Router I in the example.
Area border router	Within an AS, an area border router (ABR) connects one or more areas to the backbone. The ABR keeps a copy of the link-state database for every area it connects to. It may keep multiple copies of the link state database. An ABR summarizes learned information from one of its attached areas before it is sent to other connected areas. An ABR can connect to many areas in an AS and is considered a member of each area it connects to—shown as Router H in the example.
Autonomous system border router	The autonomous system border router (ASBR) connects to more than one AS and exchanges information with the routers in other ASs. The ASBR connects to a non-IGP such as BGP or uses static routes—shown as Router N in the example.
Internal router	The internal router (IR) has adjacencies with ONLY routers in the same area—shown as Routers E, F, I, K, and M in the example.

Designated and backup designated routers

OSPF elects a designated router (DR) and a backup designated router (BDR). The DR generates LSAs for the entire multiaccess network. Designated routers allow a reduction in network traffic and in the size of the topological database.

Designated router	Maintains a complete topology table of the network and sends updates to the other routers via multicast. All routers in an area form a slave/master relationship with the DR. Every time a router sends an update, the router sends it to the DR and BDR. The DR sends the update to all other routers in the area.

Backup designated router	Router that takes over if the DR fails.

Each router exchanges information with the DR and BDR. The DR and BDR relay information to other routers. On broadcast network segments, the number of OSPF packets reduces by the DR sending OSPF updates to a multicast IP address that all OSPF routers on the network segment are listening on.

DRs and BDRs are configurable. If you do not define the DR or BDR, OS10 assigns them per the protocol. To determine which routers are the DR and BDR, OSPF looks at the priority of the routers on the segment. The default router priority is 1. The router with the highest priority is elected DR. If there is a tie, the router with the higher router ID takes precedence. After the DR is elected, the BDR is elected the same way. A router with a router priority set to zero cannot become a DR or BDR.

Link-state advertisements

A link-state advertisement (LSA) communicates the router's routing topology to all other routers in the network.

Type 1—Router LSA	Router lists links to other routers or networks in the same area. Type 1 LSAs flood across their own area only. The link-state ID of the Type 1 LSA is the originating router ID.
Type 2—Network LSA	DR in an area lists which routers are joined within the area. Type 2 LSAs flood across their own area only. The link-state ID of the Type 2 LSA is the IP interface address of the DR.
Type 3—Summary LSA (OSPFv2), Inter-Area Prefix LSA (OSPFv3)	ABR takes information it has learned on one of its attached areas and summarizes it before sending it out on other areas it connects to. The link-state ID of the Type 3 LSA is the destination network's IP address.
Type 4—AS Border Router Summary LSA (OSPFv2), Inter-Area-Router LSA (OSPFv3)	In some cases, Type 5 External LSAs flood to areas where the detailed next-hop information may not be available because it may be using a different routing protocol. The ABR floods the information for the router, the ASBR where the Type 5 originated. The link-state ID for Type 4 LSAs is the router ID of the described ASBR.
Type 5—AS-External LSA	LSAs contain information imported into OSPF from other routing processes. Type 5 LSAs flood to all areas except stub areas. The link-state ID of the Type 5 LSA is the external network number.
Type 7—NSSA-External LSA (OSPFv2), LSA (OSPFv3)	Routers in an NSSA do not receive external LSAs from ABRs but send external routing information for redistribution. They use Type 7 LSAs to tell the ABRs about these external routes, which the ABR then translates to Type 5 external LSAs and floods as normal to the rest of the OSPF network.
Type 8—Link LSA (OSPFv3)	Type 8 LSA carries the IPv6 address information of the local links.
Type 9—Link-Local Opaque LSA (OSPFv2), Intra-Area Prefix LSA (OSPFv3)	Link-local *opaque* LSA as defined by RFC2370 for OSPFv2. Intra-Area-Prefix LSA carries the IPv6 prefixes of the router and network links for OSPFv3.
Type 11—Grace LSA (OSPFv3)	Link-local *opaque* LSA for OSPFv3 only is sent during a graceful restart by an OSPFv3 router.

The LSA header is common to LSA types. Its size is 20 bytes. One of the fields of the LSA header is the link-state ID. Each router link is defined as one of four types—type 1, 2, 3, or 4. The LSA includes a link ID field that identifies the object this link connects to, by the network number and mask. Depending on the type, the link ID has different meanings.

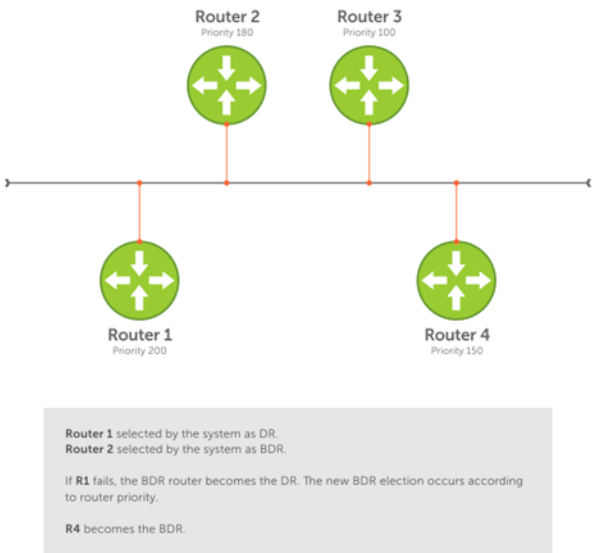
1	Point-to-point connection to another router or neighboring router

2	Connection to a transit network IP address of the DR
3	Connection to a stub network IP network or subnet number
4	Virtual link neighboring router ID

Router priority

Router priority determines the designated router for the network. The default router priority is 1. When two routers attach to a network, both attempt to become the DR. The router with the higher router priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. A router with a router priority set to zero cannot become the DR or BDR.

If not assigned, the system selects the router with the highest priority as the DR. The second highest priority is the BDR. Priority rates from 0 to 255, with 255 as the highest number with the highest priority.



OSPF route limit

OS10 supports up to 16,000 OSPF routes. Within this range, the only restriction is on intra-area routes that scale only up to 1000 routes. Other OSPF routes can scale up to 16 K.

Shortest path first throttling

Use shortest path first (SPF) throttling to delay SPF calculations during periods of network instability. In an OSPF network, a topology change event triggers an SPF calculation that is performed after a start time. When the start timer finishes, a hold time can delay the next SPF calculation for an additional time.

When the hold timer is running:

- Each time a topology change occurs, the SPF calculation delays for double the configured hold time up to maximum wait time.
- If no topology change occurs, an SPF calculation is performed and the hold timer is reset to its configured value.

Set the start, hold, and wait timers according to the stability of the OSPF network topology. Enter the values in milliseconds (ms). If you do not specify a start-time, hold-time, or max-wait value, the default values are used.

OSPFv2 and OSPFv3 instances support SPF throttling. By default, SPF timers are disabled in an OSPF instance. Enter the `no` version of this command to remove the configured SPF timers and disable SPF throttling.

1 Configure an OSPF instance from CONFIGURATION mode, from 1 to 65535.

```
router {ospf | ospfv3} instance-number
```

2 Set OSPF throttling timers in OSPF INSTANCE mode.

```
timers spf [start-time [hold-time [max-wait]]]
```

- *start-time* — Configure the initial delay before performing an SPF calculation after a topology change, from 1 to 600000 milliseconds; default 1000.
- *hold-time* — Configure the additional delay before performing an SPF calculation when a new topology change occurs, from 1 to 600000 milliseconds; default 10000.
- *max-wait* — Configure the maximum amount of hold time that can delay an SPF calculation, from 1 to 600000 milliseconds; default 10000.

Enable SPF throttling (OSPFv2)

```
OS10(config)# router ospf 100
OS10(config-router-ospf-100)# timers spf 1200 2300 3400
```

Enable SPF throttling (OSPFv3)

```
OS10(config)# router ospfv3 10
OS10(config-router-ospf-10)# timers spf 2000 3000 4000
```

View OSPFv2 SPF throttling

```
OS10(config-router-ospf-100)# do show ip ospf

Routing Process ospf 100 with ID 12.1.1.1
Supports only single TOS (TOS0) routes
It is Flooding according to RFC 2328
SPF schedule delay 1200 msecs, Hold time between two SPFs 2300 msecs
Convergence Level 0
Min LSA origination 0 msec, Min LSA arrival 1000 msec
Min LSA hold time 5000 msec, Max LSA wait time 5000 msec
Number of area in this router is 1, normal 1 stub 0 nssa 0
Area (0.0.0.1)
Number of interface in this area is 1
SPF algorithm executed 1 times
```

View OSPFv3 SPF throttling

```
OS10(config-router-ospfv3-100)# timers spf 1345 2324 9234

OS10(config-router-ospfv3-100)# do show ipv6 ospf
Routing Process ospfv3 100 with ID 129.240.244.107
SPF schedule delay 1345 msecs, Hold time between two SPFs 2324 msecs
Min LSA origination 5000 msec, Min LSA arrival 1000 msec
Min LSA hold time 0 msec, Max LSA wait time 0 msec
Number of area in this router is 1, normal 1 stub 0 nssa
Area (0.0.0.1)
Number of interface in this area is 1
SPF algorithm executed 2 times
```

OSPFv2

OSPFv2 supports IPv4 address families. OSPFv2 routers initially exchange `hello` messages to set up adjacencies with neighbor routers. The `hello` process establishes adjacencies between routers of the AS. It is not required that every router within the AS areas establish adjacencies. If two routers on the same subnet agree to become neighbors through this process, they begin to exchange network topology information in the form of LSAs.

In OSPFv2, neighbors on broadcast and non-broadcast multiple access (NBMA) network links are identified by their interface addresses, while neighbors on other types of links are identified by router-identifiers (RID).

Enable OSPFv2

OSPFv2 is disabled by default. Configure at least one interface as either Physical or Loopback and assign an IP address to the interface. You can assign any area besides area 0 a number ID. The OSPFv2 process starts automatically when you configure it globally and you can enable it for one or more interfaces.

1 Enable OSPF globally and configure an OSPF instance in CONFIGURATION mode.

```
router ospf instance-number
```

2 Enter the interface information to configure the interface for OSPF in INTERFACE mode.

```
interface ethernet node/slot/port[:subport]
```

3 Enable the interface in INTERFACE mode.

```
no shutdown
```

4 Disable the default switchport configuration and remove it from an interface or a LAG port in INTERFACE mode.

```
no switchport
```

5 Assign an IP address to the interface in INTERFACE mode.

```
ip address ip-address/mask
```

6 Enable OSPFv2 on an interface in INTERFACE mode.

```
ip ospf process-id area area-id
```

- *process-id*—Enter the OSPFv2 process ID for a specific OSPF process, from 1 to 65535.
- *area-id*—Enter the OSPFv2 area ID as an IP address (A.B.C.D) or number, from 1 to 65535.

Enable OSPFv2 configuration

```
OS10(config)# router ospf 100
OS10(conf-router-ospf-100)# exit
OS10(config)#  interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# ip address 11.1.1.1/24
OS10(conf-if-eth1/1/1)# ip ospf 100 area 0.0.0.0
```

View OSPFv2 configuration

```
OS10# show running-configuration ospf
!
interface ethernet1/1/1
 ip ospf 100 area 0.0.0.0
!
router ospf 100
...
```

Enable OSPFv2 in a non-default VRF instance

To enable OSPFv2 in a non-default VRF instance:

1 Create a non-default VRF instance in which you want to enable OSPFv2:

```
ip vrf vrf-name
```

2 Enable OSPF and configure an OSPF instance in VRF CONFIGURATION mode.

```
router ospf instance-number vrf vrf-name
```

3 Enter the interface information to configure the interface for OSPF in INTERFACE mode.

```
interface ethernet node/slot/port[:subport]
```

4 Enable the interface in INTERFACE mode.

```
no shutdown
```

5 Disable the default switchport configuration and remove it from an interface or a LAG port in INTERFACE mode.

```
no switchport
```

6 Associate the interface with the non-default VRF instance that you created earlier.

```
ip vrf forwarding vrf-name
```

7 Assign an IP address to the interface.

```
ip address ip-address/mask
```

8 Enable OSPFv2 on the interface.

```
ip ospf process-id area area-id
```

- *process-id*—Enter the OSPFv2 process ID for a specific OSPF process, from 1 to 65535.
- *area-id*—Enter the OSPFv2 area ID as an IP address (A.B.C.D) or number, from 1 to 65535.

Enable OSPFv2 configuration

```
OS10(config)# ip vrf vrf-blue
OS10(config-vrf-blue)# router ospf 100 vrf-blue
OS10(conf-router-ospf-100)# exit
OS10(config)#  interface ethernet 1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/2)# ip vrf forwarding vrf-blue
OS10(conf-if-eth1/1/1)# ip address 11.1.1.1/24
OS10(conf-if-eth1/1/1)# ip ospf 100 area 0.0.0.0
```

Assign router identifier

For managing and troubleshooting purposes, you can assign a router ID for the OSPFv2 process. Use the router's IP address as the router ID.

- Assign the router ID for the OSPFv2 process in ROUTER-OSPF mode

```
router-id ip-address
```

Assign router ID

```
OS10(config)# router ospf 10
OS10(conf-router-ospf-10)# router-id 10.10.1.5
```

View OSPFv2 status

```
OS10# show ip ospf 10
Routing Process ospf 10 with ID 10.10.1.5
Supports only single TOS (TOS0) routes
It is an Autonomous System Boundary Router
It is Flooding according to RFC 2328
Convergence Level 0
Min LSA origination 0 msec, Min LSA arrival 1000 msec
Min LSA hold time 5000 msec, Max LSA wait time 5000 msec
Number of area in this router is 1, normal 1 stub 0 nssa 0
   Area (0.0.0.0)
       Number of interface in this area is 3
       SPF algorithm executed 38 times
       Area ranges are
```

Stub areas

Type 5 LSAs are not flooded into stub areas. The ABR advertises a default route into the stub area where it is attached. Stub area routers use the default route to reach external destinations.

1 Enable OSPF routing and enter ROUTER-OSPF mode, from 1 to 65535.

```
router ospf instance number
```

2 Configure an area as a stub area in ROUTER-OSPF mode.

```
area area-id stub [no-summary]
```

- *area-id*—Enter the OSPF area ID as an IP address in A.B.C.D format or number, from 1 to 65535.
- no-summary—(Optional) Enter to prevent an ABR from sending summary LSA to the stub area.

Configure stub area

```
OS10(config)# router ospf 10
OS10(conf-router-ospf-10)# area 10.10.5.1 stub
```

View stub area configuration

```
OS10# show ip ospf
Routing Process ospf 10 with ID 130.6.196.14
Supports only single TOS (TOS0) routes
It is Flooding according to RFC 2328
SPF schedule delay 1000 msecs, Hold time between two SPFs 10000 msecs
Convergence Level 0
Min LSA origination 0 msec, Min LSA arrival 1000 msec
Min LSA hold time 5000 msec, Max LSA wait time 5000 msec
Number of area in this router is 1, normal 0 stub 1 nssa 0
   Area  (10.10.5.1)
       Number of interface in this area is 0
       SPF algorithm executed 1 times
       Area ranges are
```

```
OS10# show running-configuration ospf
!
router ospf 10
 area 10.10.5.1 stub
```

Passive interfaces

A passive interface does not send or receive routing information. Configuring an interface as a passive interface suppresses both receiving and sending routing updates.

Although the passive interface does not send or receive routing updates, the network on that interface is included in OSPF updates sent through other interfaces.

1 Enter an interface type in INTERFACE mode.

```
interface ethernet node/slot/port[:subport]
```

2 Configure the interface as a passive interface in INTERFACE mode.

```
ip ospf passive
```

Configure passive interfaces

```
OS10(config)# interface ethernet 1/1/6
OS10(conf-if-eth1/1/6)# ip ospf passive
```

View passive interfaces

```
OS10# show running-configuration
!!!
!!
interface ethernet1/1/6
 ip address 10.10.10.1/24
 no switchport
 no shutdown
 ip ospf 100 area 0.0.0.0
 ip ospf passive
```

```
!!
!
```

You can disable a passive interface using the `no ip ospf passive` command.

Fast convergence

Fast convergence sets the minimum origination and arrival LSA parameters to zero (0), allowing rapid route calculation. A higher convergence level can result in occasional loss of OSPF adjacency.

Convergence level 1 meets most convergence requirements. The higher the number, the faster the convergence, and the more frequent the route calculations and updates. This impacts CPU utilization and may impact adjacency stability in larger topologies.

NOTE: Select higher convergence levels only after checking with Dell EMC Technical Support.

When you disable fast-convergence, origination and arrival LSA parameters are set to 0 msec and 1000 msec, respectively. Setting the convergence parameter from 1 to 4 indicates the actual convergence level. Each convergence setting adjusts the LSA parameters to zero, but the `convergence-level` parameter changes the convergence speed. The higher the number, the faster the convergence.

- Enable OSPFv2 fast-convergence and enter the convergence level in ROUTER-OSPF mode, from 1 to 4.
  `fast-converge` *`convergence-level`*

Configure fast convergence

```
OS10(config)# router ospf 65535
OS10(conf-router-ospf-65535)# fast-converge 1
```

View fast convergence

```
OS10(conf-router-ospf-65535)# do show ip ospf

Routing Process ospf 65535 with ID 99.99.99.99
Supports only single TOS (TOS0) routes
It is an Autonomous System Border Router
It is an Area Border Router
It is Flooding according to RFC 2328
Convergence Level 1
Min LSA origination 0 msec, Min LSA arrival 0 msec
Min LSA hold time 0 msec, Max LSA wait time 0 msec
Number of area in this router is 3, normal 1 stub 1 nssa 1
   Area BACKBONE (0)
       Number of interface in this area is 1
       SPF algorithm executed 28 times
       Area ranges are

   Area  (2)
       Number of interface in this area is 1
       SPF algorithm executed 28 times
       Area ranges are

   Area  (3)
       Number of interface in this area is 1
       SPF algorithm executed 28 times
       Area ranges are
```

Disable fast convergence

```
OS10(conf-router-ospf-65535)# no fast-converge
```

Interface parameters

To avoid routing errors, interface parameter values must be consistent across all interfaces. For example, set the same time interval for the hello packets on all routers in the OSPF network to prevent misconfiguration of OSPF neighbors.

1 To change the OSPFv2 parameters in CONFIGURATION mode, enter the interface.

```
interface interface-name
```

2 Change the cost associated with OSPF traffic on the interface in INTERFACE mode, from 1 to 65535. The default depends on the interface speed.

```
ip ospf cost
```

3 Change the time interval, from 1 to 65535, that the router waits before declaring a neighbor dead in INTERFACE mode. The default time interval is 40. The dead interval must be four times the hello interval and must be the same on all routers in the OSPF network.

```
ip ospf dead-interval seconds
```

4 Change the time interval between hello-packet transmission in INTERFACE mode, from 1 to 65535. The default time interval is 10. The hello interval must be the same on all routers in the OSPF network.

```
ip ospf hello-interval seconds
```

5 Change the priority of the interface, which determines the DR for the OSPF broadcast network in INTERFACE mode, from 0 to 255. The default priority of the interface is 1.

```
ip ospf priority number
```

6 Change the retransmission interval time, in seconds, between LSAs in INTERFACE mode, from 1 to 3600. The default retransmission interval time is 5. The retransmit interval must be the same on all routers in the OSPF network.

```
ip ospf retransmit-interval seconds
```

7 Change the wait period between link state update packets sent out the interface in INTERFACE mode, from 1 to 3600. The default wait period is 1. The transmit delay must be the same on all routers in the OSPF network.

```
ip ospf transmit-delay seconds
```

Change parameters and view interface status

```
OS10(conf-if-eth1/1/1)# ip ospf hello-interval 5
OS10(conf-if-eth1/1/1)# ip ospf dead-interval 20
OS10(conf-if-eth1/1/1)# ip ospf retransmit-interval 30
OS10(conf-if-eth1/1/1)# ip ospf transmit-delay 200
```

View OSPF interface configuration

```
OS10(conf-if-eth1/1/1)# do show ip ospf interface

ethernet1/1/1 is up, line protocol is up
  Internet Address 11.1.1.1/24, Area 0.0.0.0
  Process ID 65535, Router ID 99.99.99.99, Network Type broadcast, Cost: 1
  Transmit Delay is 200 sec, State BDR, Priority 1
  Designated Router (ID) 150.1.1.1, Interface address 11.1.1.2
  Backup Designated router (ID) 99.99.99.99, Interface address 11.1.1.1
  Timer intervals configured, Hello 5, Dead 20, Wait 20, Retransmit 30
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 150.1.1.1(Designated Router)
```

Redistribute routes

Add routes from other routing instances or protocols to the OSPFv2 process and include BGP, static, or connected routes in the OSPFv2 process. Do not route IBGP routes to OSPFv2 unless there are route-maps associated with the OSPFv2 redistribution.

- Enter which routes redistribute into the OSPFv2 process in ROUTER-OSPF mode.

  ```
  redistribute {bgp as-number| connected | static} [route-map map-name]
  ```

  – `bgp | connected | static`—Enter a keyword to redistribute those routes.

– `route-map map-name`—Enter the name of a configured route map.

Configure redistribute routes

```
OS10(conf-router-ospf-10)# redistribute bgp 4 route-map aloha
OS10(conf-router-ospf-10)# redistribute connected route-map aloha
OS10(conf-router-ospf-10)# redistribute static route-map aloha
```

View OSPF configuration

```
OS10(conf-router-ospf-10)# do show running-configuration ospf
!
router ospf 10
 redistribute bgp 4 route-map aloha
 redistribute connected route-map aloha
 redistribute static route-map aloha
!
```

Default route

You can generate an external default route and distribute the default information to the OSPFv2 routing domain.

- Generate the default route using the `default-information originate [always]` command in ROUTER-OSPF mode.

Configure default route

```
OS10(config)# router ospf 10
OS10(config-router-ospf-10)# default-information originate always
```

View default route configuration

```
OS10(config-router-ospf-10)# show configuration
!
router ospf 10
 default-information originate always
```

Summary address

You can configure a summary address for an ASBR to advertise one external route as an aggregate, for all redistributed routes that are covered by specified address range.

- Configure the summary address in ROUTER-OSPF mode.

  `summary-address ip-address/mask [not-advertise | tag tag-value]`

Configure summary address

```
OS10(config)# router ospf 100
OS10(config-router-ospf-100)# summary-address 10.0.0.0/8 not-advertise
```

View summary address

```
OS10(config-router-ospf-100)# show configuration
!
router ospf 100
  summary-address 10.0.0.0/8 not-advertise
```

Graceful restart

When a networking device restarts, the adjacent neighbors and peers detect the condition. During a graceful restart, the restarting device and neighbors continue to forward the packets without interrupting network performance. The neighbors that help in the restart process are called helper routers.

When you enable graceful restart, the restarting device retains the routes learned by OSPF in the forwarding table. To re-establish OSPF adjacencies with neighbors, the restart OSPF process sends a grace LSA to all neighbors. In response, the helper router enters Helper mode and sends an acknowledgement back to the restarting device.

OS10 supports graceful restart Helper mode. Use the `graceful-restart role helper-only` command to enable Helper mode in ROUTER OSPF mode.

```
OS10(config)# router ospf 10
OS10(conf-router-ospf-10)# graceful-restart role helper-only
```

Use the `no` version of the command to disable Helper mode.

OSPFv2 authentication

You can enable OSPF authentication either with clear text or MD5.

- Set a clear text authentication scheme on the interface in INTERFACE mode.
  `ip ospf authentication-key key`
- Set MD5 authentication in INTERFACE mode.
  `ip ospf message-digest-key keyid md5 key`

Configure text authentication

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ip ospf authentication-key sample
```

View text authentication

```
OS10(conf-if-eth1/1/1)# show configuration
!
interface ethernet1/1/1
 ip address 10.10.10.2/24
 no switchport
 no shutdown
 ip ospf 100 area 0.0.0.0
 ip ospf authentication-key sample
```

Configure MD5 authentication

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ip ospf message-digest-key 2 md5 sample12345
```

View MD5 authentication

```
OS10(conf-if-eth1/1/1)# show configuration

!
interface ethernet1/1/1
 ip address 10.10.10.2/24
 no switchport
 no shutdown
 ip ospf 100 area 0.0.0.0
 ip ospf message-digest-key 2 md5 sample12345
```

Troubleshoot OSPFv2

You can troubleshoot OSPFv2 operations, and check questions for typical issues that interrupt a process.

- Is OSPF enabled globally?
- Is OSPF enabled on the interface?
- Are adjacencies established correctly?
- Are the interfaces configured for L3 correctly?

- Is the router in the correct area type?
- Are the OSPF routes included in the OSPF database?
- Are the OSPF routes included in the routing table in addition to the OSPF database?
- Are you able to ping the IPv4 address of adjacent router interface?

Troubleshooting OSPF with show commands

- View a summary of all OSPF process IDs enabled in EXEC mode.
  ```
  show running-configuration ospf
  ```
- View summary information of IP routes in EXEC mode.
  ```
  show ip route summary
  ```
- View summary information for the OSPF database in EXEC mode.
  ```
  show ip ospf database
  ```
- View the configuration of OSPF neighbors connected to the local router in EXEC mode.
  ```
  show ip ospf neighbor
  ```
- View routes that OSPF calculates in EXEC mode.
  ```
  show ip ospf routes
  ```

View OSPF configuration

```
OS10# show running-configuration ospf
!
interface ethernet1/1/1
ip ospf 100 area 0.0.0.0
!
router ospf 100
log-adjacency-changes
```

OSPFv2 commands

area default-cost

Sets the metric for the summary default route generated by the ABR and sends it to the stub area.

Syntax	`area` *`area-id`* `default-cost` *`cost`*
Parameters	• *`area-id`* — Enter the OSPF area in dotted decimal A.B.C.D format or enter a number, from 0 to 65535. • *`cost`* — Enter a cost for the stub area's advertised external route metric, from 0 to 65535.
Default	Cost is 1
Command Mode	ROUTER-OSPF
Usage Information	The cost is also referred as *reference-bandwidth* or *bandwidth*. Use the `area default-cost` command on the border routers at the edge of a stub area. The `no` version of this command resets the value to the default.
Example	`OS10(conf-router-ospf-10)# area 10.10.1.5 default-cost 10`
Supported Releases	10.2.0E or later

area nssa

Defines an area as a NSSA.

Syntax

```
area area-id nssa [default-information-originate | no-redistribution | no-summary]
```

Parameters

- *area-id* — Enter the OSPF area ID as an IP address in A.B.C.D format or number, from 1 to 65535.
- `no-redistribution` — (Optional) Prevents the `redistribute` command from distributing routes into the NSSA. Use `no-redistribution` command only in an NSSA ABR.
- `no-summary` — (Optional) Ensures that no summary LSAs are sent to the NSSA.

Default Not configured

Command Mode ROUTER-OSPF

Usage Information The `no` version of this command deletes an NSSA.

Example

```
OS10(conf-router-ospf-10)# area 10.10.1.5 nssa
```

Supported Releases 10.2.0E or later

area range

Summarizes routes matching an address/mask at an area in ABRs.

Syntax

```
area area-id range ip-address [no-advertise]
```

Parameters

- *area-id* — Set the OSPF area ID as an IP address in A.B.C.D format or number, from 1 to 65535.
- *ip-address* — (Optional) Enter an IP address/mask in dotted decimal format.
- `no-advertise` — (Optional) Set the status to *Do Not Advertise*. The Type 3 summary-LSA is suppressed and the component networks remain hidden from other areas.

Default Not configured

Command Mode ROUTER-OSPF

Usage Information The `no` version of this command disables the route summarizations.

Example

```
OS10(conf-router-ospf-10)# area 0 range 10.1.1.4/8 no-advertise
```

Supported Releases 10.2.0E or later

area stub

Defines an area as the OSPF stub area.

Syntax

```
area area-id stub [no-summary]
```

Parameters

- *area-id*—Set the OSPF area ID as an IP address in A.B.C.D format or number, from 1 to 65535.
- `no-summary`—(Optional) Prevents an ABR from sending summary LAs into the stub area.

Default Not configured

Command Mode ROUTER-OSPF

Usage Information	The `no` version of this command deletes a stub area.
Example	`OS10(config)# router ospf 10` `OS10(conf-router-ospf-10)# area 10.10.1.5 stub`
Supported Releases	10.2.0E or later

auto-cost reference-bandwidth

Calculates default metrics for the interface based on the configured auto-cost reference bandwidth value.

Syntax	`auto-cost reference-bandwidth value`
Parameters	*value* — Enter the reference bandwidth value to calculate the OSPF interface cost in megabits per second, from 1 to 4294967.
Default	100000
Command Mode	ROUTER-OSPF
Usage Information	The value set by the `ip ospf cost` command in INTERFACE mode overrides the cost resulting from the `auto-cost` command. The `no` version of this command resets the value to the default.
Example	`OS10(config)# router ospf 10` `OS10(conf-router-ospf-10)# auto-cost reference-bandwidth 150`
Supported Releases	10.2.0E or later

clear ip ospf process

Clears all OSPF routing tables.

Syntax	`clear ip ospf {instance-number} [vrf vrf-name] process`
Parameters	• *instance-number* — Enter an OSPF instance number, from 1 to 65535. • `vrf vrf-name` — Enter the keyword `vrf` followed by the name of the VRF to reset the OSPF process configured in that VRF.
Default	Not configured
Command Mode	EXEC
Usage Information	This command clears all entries in the OSPF routing table.
Example	`OS10# clear ip ospf 3 vrf vrf-test process`
Supported Releases	10.2.0E or later

clear ip ospf statistics

Clears OSPF traffic statistics.

Syntax	`clear ip ospf [instance-number] [vrf vrf-name] statistics`
Parameters	• *instance-number* — (Optional) Enter an OSPF instance number, from 1 to 65535. • `vrf vrf-name` — (Optional) Enter the keyword `vrf` followed by the name of the VRF to clear OSPF traffic statistics in that configured VRF.

Default	Not configured
Command Mode	EXEC
Usage Information	This command clears the OSPF traffic statistics in a specified instance or in all the configured OSPF instances, and resets them to zero.
Example	`OS10# clear ip ospf 10 vrf vrf-test statistics`
Supported Releases	10.4.0E(R1) or later

default-information originate

Generates and distributes a default external route information to the OSPF routing domain.

Syntax	`default-information originate [always]`
Parameters	`always` — (Optional) Always advertise the default route.
Defaults	Disabled
Command Mode	ROUTER-OSPF
Usage Information	The `no` version of this command disables the distribution of default route.
Example	`OS10(config)# router ospf 10` `OS10(config-router-ospf-10)# default-information originate always`
Supported Releases	10.3.0E or later

default-metric

Assigns a metric value to redistributed routes for the OSPF process.

Syntax	`default-metric` *`number`*
Parameters	*`number`* — Enter a default-metric value, from 1 to 16777214.
Default	Not configured
Command Mode	ROUTER-OSPF
Usage Information	The `no` version of this command disables the default-metric configuration.
Example	`OS10(conf-router-ospf-10)# default-metric 2000`
Supported Releases	10.2.0E or later

fast-converge

Sets the minimum LSA origination and arrival times to zero (0) allowing more rapid route computation so convergence takes less time.

Syntax	`fast-converge` *`convergence-level`*
Parameters	*`convergence-level`* — Enter a desired convergence level value, from 1 to 4.
Default	Not configured
Command Mode	ROUTER-OSPF
Usage Information	Convergence level 1 (optimal) meets most convergence requirements.

NOTE: Only select higher convergence levels following consultation with Dell EMC Technical Support.

The `no` version of this command disables the fast-convergence configuration.

Example

```
OS10(conf-router-ospf-10)# fast-converge 3
```

Supported Releases 10.2.0E or later

graceful-restart

Enables Helper mode during a graceful or hitless restart.

Syntax `graceful-restart role helper-only`

Parameters None

Defaults Disabled

Command Mode ROUTER-OSPF

Usage Information The `no` version of this command disables Helper mode.

Example

```
OS10(config)# router ospf 10
OS10(conf-router-ospf-10)# graceful-restart role helper-only
```

Supported Releases 10.3.0E or later

ip ospf area

Attaches an interface to an OSPF area.

Syntax `ip ospf process-id area area-id`

Parameters

- *process-id* — Set an OSPF process ID for a specific OSPF process, from 1 to 65535.
- `area` *area-id* — Enter the OSPF area ID in dotted decimal A.B.C.D format or enter an area ID number, from 1 to 65535.

Default Not configured

Command Mode INTERFACE

Usage Information The `no` version of this command removes an interface from an OSPF area.

Example

```
OS10(conf-if-vl-10)# ip ospf 10 area 5
```

Supported Releases 10.2.0E or later

ip ospf authentication-key

Configures a text authentication key to enable OSPF traffic on an interface.

Syntax `ip ospf authentication-key key`

Parameters *key* — Enter an eight-character string for the authentication key.

Defaults Not configured

Command Mode INTERFACE

Usage Information	To exchange OSPF information, all neighboring routers in the same network must use the same authentication key. The `no` version of this command deletes the authentication key.
Example	`OS10(config)# interface ethernet 1/1/1` `OS10(conf-if-eth1/1/1)# ip ospf authentication-key sample`
Supported Releases	10.3.0E or later

ip ospf cost

Changes the cost associated with the OSPF traffic on an interface.

Syntax	`ip ospf cost` *`cost`*
Parameters	*`cost`* — Enter a value as the OSPF cost for the interface, from 1 to 65335.
Default	Based on bandwidth reference
Command Mode	INTERFACE
Usage Information	if not configured, interface cost is based on the `auto-cost` command. This command configures OSPF over multiple vendors to ensure that all routers use the same cost. If you manually configure the cost, the calculated cost based on the reference bandwidth does not apply to the interface. The `no` version of this command removes the IP OSPF cost configuration.
Example	`OS10(config)# interface vlan 10` `OS10(conf-if-vl-1)# ip ospf cost 10`
Supported Releases	10.2.0E or later

ip ospf dead-interval

Sets the time interval since the last hello-packet was received from a router. After the interval elapses, the neighboring routers declare the router dead.

Syntax	`ip ospf dead-interval` *`seconds`*
Parameters	*`seconds`* — Enter the dead interval value in seconds, from 1 to 65535.
Default	40 seconds
Command Mode	INTERFACE
Usage Information	The dead interval is four times the default hello-interval by default. The `no` version of this command resets the value to the default.
Example	`OS10(conf-if-vl-10)# ip ospf dead-interval 10`
Supported Releases	10.2.0E or later

ip ospf hello-interval

Sets the time interval between the hello packets sent on the interface.

Syntax	`ip ospf hello-interval` *`seconds`*
Parameters	*`seconds`* — Enter the hello-interval value in seconds, from 1 to 65535.
Default	10 seconds
Command Mode	INTERFACE

Usage Information	All routers in a network must have the same hello time interval between the hello packets. The `no` version of the this command resets the value to the default.
Example	`OS10(conf-if-vl-10)# ip ospf hello-interval 30`
Supported Releases	10.2.0E or later

ip ospf message-digest-key

Enables OSPF MD5 authentication and sends an OSPF message digest key on the interface.

Syntax	`ip ospf message-digest-key` *keyid* `md5` *key*
Parameters	• *keyid* — Enter an MD5 key ID for the interface, from 1 to 255. • *key* — Enter a character string as the password. A maximum of 16 characters.
Defaults	Not configured
Command Mode	INTERFACE
Usage Information	All neighboring routers in the same network must use the same key value to exchange OSPF information. The `no` version of this command deletes the authentication key.
Example	`OS10(config)# interface ethernet 1/1/1` `OS10(conf-if-eth1/1/1)# ip ospf message-digest-key 2 md5 sample12345`
Supported Releases	10.3.0E or later

ip ospf mtu-ignore

Enables OSPF MTU mismatch detection on receipt of DBD packets.

Syntax	`ip ospf mtu-ignore`
Parameters	None
Default	Not configured
Command Mode	INTERFACE
Usage Information	When neighbors exchange DBD packets, the OSPF process checks if the neighbors are using the same MTU on a common interface. If the receiving MTU in the DBD packet is higher than the IP MTU configured on the incoming interface, OSPF adjacency does not establish. The `no` version of this command disables the IP OSPF mtu-ignore configuration.
Example	`OS10(conf-if-vl-10)# ip ospf mtu-ignore`
Supported Releases	10.2.0E or later

ip ospf network

Sets the network type for the interface.

Syntax	`ip ospf network {point-to-point | broadcast}`
Parameters	• `point-to-point` — Sets the interface as part of a point-to-point network. • `broadcast` — Sets the interface as part of a broadcast network.

Default	Broadcast
Command Mode	INTERFACE
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(conf-if-eth1/1/1)# ip ospf network broadcast`
Supported Releases	10.2.0E or later

ip ospf passive

Configures an interface as a passive interface and suppresses both receiving and sending routing updates to the passive interface.

Syntax	`ip ospf passive`
Parameters	None
Default	Not configured
Command Mode	INTERFACE
Usage Information	You must configure the interface before setting the interface to Passive mode. The `no` version of the this command disables the passive interface configuration. **NOTE: As loopback interfaces are implicitly passive, the configuration to suppress sending and receiving of OSPF routing updates does not take effect on the loopback interfaces. However, network information corresponding to these loopback interfaces is still announced in OSPF LSAs that are sent through other interfaces configured for OSPF.**
Example	`OS10(conf-if-eth1/1/6)# ip ospf passive`
Supported Releases	10.2.0E or later

ip ospf priority

Sets the priority of the interface to determine the DR for the OSPF network.

Syntax	`ip ospf priority` *number*
Parameters	*number* — Enter a router priority number, from 0 to 255.
Default	1
Command Mode	INTERFACE
Usage Information	When two routers attached to a network attempt to become the DR, the one with the higher router priority takes precedence. The `no` version of this command resets the value to the default.
Example	`OS10(conf-if-eth1/1/6)# ip ospf priority 4`
Supported Releases	10.2.0E or later

ip ospf retransmit-interval

Sets the retransmission time between lost LSAs for adjacencies belonging to the interface.

Syntax	`ip ospf retransmit-interval` *seconds*
Parameters	*seconds* — Enter a value in seconds as the interval between retransmission, from 1 to 3600.

Default	5 seconds
Command Mode	INTERFACE
Usage Information	Set the time interval to a number large enough to avoid unnecessary retransmission. The `no` version of this command resets the value to the default.
Example	`OS10(conf-if-eth1/1/6)# ip ospf retransmit-interval 20`
Supported Releases	10.2.0E or later

ip ospf transmit-delay

Sets the estimated time required to send a link state update packet on the interface.

Syntax	`ip ospf transmit-delay` *seconds*
Parameters	*seconds* — Set the time in seconds required to send a link-state update, from 1 to 3600.
Default	1 second
Command Mode	INTERFACE
Usage Information	When you set the `ip ospf transmit-delay` value, take into account the transmission and propagation delays for the interface. The `no` version of this command resets the value to the default.
Example	`OS10(conf-if-eth1/1/4)# ip ospf transmit-delay 5`
Supported Releases	10.2.0E or later

log-adjacency-changes

Enables logging of syslog messages regarding changes in the OSPF adjacency state.

Syntax	`log-adjacency-changes`
Parameters	None
Default	Disabled
Command Mode	ROUTER-OSPF
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(config)# router ospf 10` `OS10(conf-router-ospf-10)# log-adjacency-changes`
Supported Releases	10.2.0E or later

max-metric router-lsa

Configures OSPF to advertise a maximum metric on a router so that it is not desired as an intermediate hop from other routers.

Syntax	`max-metric router-lsa`
Parameters	None
Default	Not configured
Command Mode	ROUTER-OSPF

Usage Information	Routers in the network do not prefer other routers as the next intermediate hop after they calculate the shortest path. The `no` version of this command disables the maximum metric advertisement configuration.
Example	`OS10(conf-router-ospf-10)# max-metric router-lsa`
Supported Releases	10.2.0E or later

maximum-paths

Enables forwarding of packets over multiple paths.

Syntax	`maximum–paths` *`number`*
Parameters	*`number`* —Enter the number of paths for OSPF, from 1 to 128.
Default	64
Command Mode	ROUTER-OSPF
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(config)# router ospf 10` `OS10(conf-router-ospf-10)# maximum-paths 1`
Supported Releases	10.2.0E or later

redistribute

Redistributes information from another routing protocol or routing instance to the OSPFv2 process.

Syntax	`redistribute {bgp` *`as-number`*`| connected | static} [route-map` *`map-name`*`]`
Parameters	• *`as-number`* — Enter an autonomous number to redistribute BGP routing information throughout the OSPF instance, from 1 to 4294967295. • `connected` — Enter the information from the connected active routes on interfaces to redistribute. • `static` — Enter the information from static routes on interfaces to redistribute. • *`route-map name`* — Enter the name of a configured route-map.
Defaults	Not configured
Command Mode	ROUTER-OSPF
Usage Information	When an OSPF redistributes, the process does not completely remove from the BGP configuration. The `no` version of this command disables the redistribute configuration.
Example	`OS10(config)# router ospf 10` `OS10(conf-router-ospf-10)# redistribute bgp 4 route-map dell1`
Example (Connected)	`OS10(config)# router ospf 10` `OS10(conf-router-ospf-10)# redistribute connected route-map dell2`
Supported Releases	10.2.0E or later

router-id

Configures a fixed router ID for the OSPF process.

Syntax	`router-id` *`ip-address`*

Parameters	*ip-address* — Enter the IP address of the router as the router ID.
Default	Not configured
Command Mode	ROUTER-OSPF
Usage Information	Configure an arbitrary value in the IP address format for each router. Each router ID must be unique. Use the fixed router ID for the active OSPF router process. Changing the router ID brings down the existing OSPF adjacency. The new router ID becomes effective immediately. The `no` version of this command disables the router ID configuration.
Example	`OS10(config)# router ospf 10` `OS10(conf-router-ospf-10)# router-id 10.10.1.5`
Supported Releases	10.2.0E or later

router ospf

Enters Router OSPF mode and configures an OSPF instance.

Syntax	`router ospf` *instance-number* `[vrf` *vrf-name*`]`
Parameters	• *instance-number*—Enter a router OSPF instance number, from 1 to 65535. • `vrf` *vrf-name* — Enter the keyword `vrf` followed by the name of the VRF to configure an OSPF instance in that VRF.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	Assign an IP address to an interface before using this command. The `no` version of this command deletes an OSPF instance.
Example	`OS10(config)# router ospf 10 vrf vrf-test`
Supported Releases	10.2.0E or later

show ip ospf

Displays OSPF instance configuration information.

Syntax	`show ip ospf [`*instance-number*`] [vrf` *vrf-name*`]`
Parameters	• *instance-number* — View OSPF information for a specified instance number from, 1 to 65535. • `vrf` *vrf-name* — Enter the keyword `vrf` followed by the name of the VRF to display OSPF configuration information corresponding to that VRF.
Default	Not configured
Command Mode	EXEC
Usage Information	None

Example	`OS10# show ip ospf 10`
	`Routing Process ospf 10 with ID 111.2.1.1`
	`Supports only single TOS (TOS0) routes`
	`It is an Autonomous System Boundary Router`
	`It is Flooding according to RFC 2328`
	`Convergence Level 0`
	`Min LSA origination 0 msec, Min LSA arrival 1000 msec`

```
Min LSA hold time 5000 msec, Max LSA wait time 5000 msec
Number of area in this router is 1, normal 1 stub 0 nssa 0
   Area (0.0.0.0)
       Number of interface in this area is 3
       SPF algorithm executed 38 times
       Area ranges are
```

Supported Releases 10.2.0E or later

show ip ospf asbr

Displays all the ASBR visible to OSPF.

Syntax `show ip ospf [process-id] [vrf vrf-name] asbr`

Parameters

- *process-id*—(Optional) Displays information based on the process ID.
- `vrf` *vrf-name* — (Optional) Displays the ASBR router visible to the OSPF process configured in the specified VRF.

Default Not configured

Command Mode EXEC

Usage Information You can isolate problems with external routes. External OSPF routes are calculated by adding the LSA cost to the cost of reaching the ASBR router. If an external route does not have the correct cost, this command determines if the path to the originating router is correct. ASBRs that are not in directly connected areas display. You can determine if an ASBR is in a directly connected area by the flags. For ASBRs in a directly connected area, E flags are set.

Example

```
OS10# show ip ospf 10 asbr

RouterID          Flags    Cost   Nexthop          Interface    Area
112.2.1.1         E/-/-/   1      110.1.1.2        vlan3050    0.0.0.0
111.2.1.1         E/-/-/   0      0.0.0.0          -             -
```

Supported Releases 10.2.0E or later

show ip ospf database

Displays all LSA information. You must enable OSPF to generate output.

Syntax `show ip ospf [process-id] [vrf vrf-name] database`

Parameters

- *process-id* — (Optional) View LSA information for a specific OSPF process ID. If you do not enter a process ID, the command applies to all the configured OSPF processes.
- `vrf` *vrf-name* — (Optional) Enter the keyword `vrf` followed by the name of the VRF to display LSA information for the OSPF process corresponding to that VRF.

Default Not configured

Command Mode EXEC

Usage Information

- `Link ID` — Identifies the router ID.
- `ADV Router` — Identifies the advertising router's ID.
- `Age` — Displays the LS age.
- `Seq#` — Identifies the LS sequence number. This identifies old or duplicate LSAs.
- `Checksum` — Displays the Fletcher checksum of an LSA's complete contents.

- `Link count` — Displays the number of interfaces for that router.

Example

```
OS10# show ip ospf 10 database
OSPF Router with ID (111.2.1.1) (Process ID 10)

                Router (Area 0.0.0.0)

Link ID         ADV Router        Age     Seq#          Checksum    Link count
111.2.1.1       111.2.1.1         1281    0x8000000d    0x9bf2          3
111.111.111.1   111.111.111.1     1430    0x8000021a    0x515a          1
111.111.111.2   111.111.111.2     1430    0x8000021a    0x5552          1
112.2.1.1       112.2.1.1         1282    0x8000000b    0x0485          3
112.112.112.1   112.112.112.1     1305    0x80000250    0xbab2          1
112.112.112.2   112.112.112.2     1305    0x80000250    0xbeaa          1

                Network (Area 0.0.0.0)

Link ID         ADV Router        Age     Seq#          Checksum
110.1.1.2       112.2.1.1         1287    0x80000008    0xd2b1
111.1.1.1       111.2.1.1         1458    0x80000008    0x1b8f
111.2.1.1       111.2.1.1         1458    0x80000008    0x198f
112.1.1.1       112.2.1.1         1372    0x80000008    0x287c
112.2.1.1       112.2.1.1         1372    0x80000008    0x267c

                Summary Network (Area 0.0.0.0)
```

Supported Releases 10.2.0E or later

show ip ospf database asbr-summary

Displays information about AS boundary LSAs.

Syntax `show ip ospf [process-id] database asbr-summary`

Parameters

- *process-id*—(Optional) Displays the AS boundary LSA information for a specified OSPF process ID. If you do not enter a process ID, this applies only to the first OSPF process.
- `vrf vrf-name` — (Optional) Displays the AS boundary LSA information for a OSPF process ID corresponding to the specified VRF.

Default Not configured

Command Mode EXEC

Usage Information

- `LS Age`—Displays the LS age.
- `Options`—Displays optional capabilities.
- `LS Type`—Displays the LS type.
- `Link State ID`—Identifies the router ID.
- `Advertising Router`—Identifies the advertising router's ID.
- `LS Seq Number`—Identifies the LS sequence number. This identifies old or duplicate LSAs.
- `Checksum`—Displays the Fletcher checksum of an LSA's complete contents.
- `Length`—Displays the LSA length in bytes.
- `Network Mask`—Identifies the network mask implemented on the area.
- `TOS`—Displays the ToS options. The only option available is zero.
- `Metric`—Displays the LSA metric.

Example

```
OS10# show ip ospf 10 database asbr-summary
```

```
  OSPF Router with ID (1.1.1.1) (Process ID 100)

   Summary Asbr (Area 0.0.0.1)

LS age: 32
Options: (No TOS-Capability, No DC)
LS type: Summary Asbr
Link State ID: 8.1.1.1
Advertising Router: 2.2.2.2
LS Seq Number: 0x80000001
Checksum: 0xB595
Length: 28
Network Mask: /0
    TOS: 0 Metric: 0
```

Supported Releases 10.2.0E or later

show ip ospf database external

Displays information about the AS external Type 5 LSAs.

Syntax `show ip ospf [process-id] [vrf vrf-name] database external`

Parameters

- *process-id*—(Optional) Displays AS external Type 5 LSA information for a specified OSPF process ID. If you do not enter a process ID, this command applies only to the first OSPF process.
- `vrf` *vrf-name* — (Optional) Displays AS external (Type 5) LSA information for a specified OSPF Process ID corresponding to a VRF.

Default Not configured

Command Mode EXEC

Usage Information

- `LS Age` — Displays the LS age.
- `Options` — Displays the optional capabilities available on the router.
- `LS Type` — Displays the LS type.
- `Link State ID` — Identifies the router ID.
- `Advertising Router` — Identifies the advertising router's ID.
- `LS Seq Number` — Identifies the LS sequence number. This identifies old or duplicate LSAs.
- `Checksum` — Displays the Fletcher checksum of an LSA's complete contents.
- `Length` — Displays the LSA length in bytes.
- `Network Mask` — Identifies the network mask implemented on the area.
- `TOS` — Displays the ToS options. The only option available is zero.
- `Metric` — Displays the LSA metric.

Example

```
OS10# show ip ospf 10 database external

 OSPF Router with ID (111.2.1.1) (Process ID 10)

                Type-5 AS External

  LS age: 1424
  Options: (No TOS-capability, No DC, E)
  LS type: Type-5 AS External
  Link State ID: 110.1.1.0
  Advertising Router: 111.2.1.1
  LS Seq Number: 0x80000009
  Checksum: 0xc69a
  Length: 36
```

```
  Network Mask: /24
      Metric Type: 2
      TOS: 0
      Metric: 20
      Forward Address: 110.1.1.1
      External Route Tag: 0
```

Supported Releases 10.2.0E or later

show ip ospf database network

Displays information about network Type 2 LSA information.

Syntax `show ip ospf [process-id] [vrf vrf-name] database network`

Parameters

- *process-id* — (Optional) Displays network Type2 LSA information for a specified OSPF process ID. If you do not enter a process ID, this command applies only to the first OSPF process.
- `vrf` *vrf-name* — (Optional) Displays network Type2 LSA information for a specified OSPF process ID corresponding to a VRF.

Default Not configured

Command Mode EXEC

Usage Information

- `LS Age`—Displays the LS age.
- `Options`—Displays optional capabilities.
- `LS Type`—Displays the LS type.
- `Link State ID`—Identifies the router ID.
- `Advertising Router`—Identifies the advertising router's ID.
- `LS Seq Number`—Identifies the LS sequence number. This identifies old or duplicate LSAs.
- `Checksum`—Displays the Fletcher checksum of an LSA's complete contents.
- `Length`—Displays the LSA length in bytes.
- `Network Mask`—Identifies the network mask implemented on the area.
- `TOS`—Displays the ToS options. The only option available is zero..
- `Metric`—Displays the LSA metric.

Example

```
OS10# show ip ospf 10 database network
OSPF Router with ID (111.2.1.1) (Process ID 10)

                Network (Area 0.0.0.0)

  LS age: 1356
  Options: (No TOS-capability, No DC, E)
  LS type: Network
  Link State ID: 110.1.1.2
  Advertising Router: 112.2.1.1
  LS Seq Number: 0x80000008
  Checksum: 0xd2b1
  Length: 32
  Network Mask: /24
      Attached Router: 111.2.1.1
      Attached Router: 112.2.1.1
```

Supported Releases 10.2.0E or later

show ip ospf database nssa external

Displays information about the NSSA-External Type 7 LSA.

Syntax

```
show ip ospf [process-id] [vrf vrf-name] database nssa external
```

Parameters

- *process-id* — (Optional) Displays NSSA-External Type7 LSA information for a specified OSPF process ID. If you do not enter a process ID, this command applies only to the first OSPF process.
- `vrf vrf-name` — (Optional) Displays NSSA-External Type7 LSA information for a specified OSPF process ID corresponding to a VRF.

Default

Not configured

Command Mode

EXEC

Usage Information

- `LS Age` — Displays the LS age.
- `Options` — Displays the optional capabilities available on the router.
- `LS Type` — Displays the LS type.
- `Link State ID` — Identifies the router ID.
- `Advertising Router` — Identifies the advertising router's ID.
- `LS Seq Number` — Identifies the LS sequence number. This identifies old or duplicate LSAs.
- `Checksum` — Displays the Fletcher checksum of an LSA's complete contents.
- `Length` — Displays the LSA length in bytes.
- `Network Mask`—Identifies the network mask implemented on the area.
- `TOS`—Displays the ToS options. The only option available is zero.
- `Metric`—Displays the LSA metric.

Example

```
OS10# show ip ospf database nssa external

       OSPF Router with ID (2.2.2.2) (Process ID 100)

          NSSA External (Area 0.0.0.1)

LS age: 98
Options: (No TOS-Capability, No DC, No Type 7/5 translation)
LS type: NSSA External
Link State ID: 0.0.0.0
Advertising Router: 1.1.1.1
LS Seq Number: 0x80000001
Checksum: 0x430C
Length: 36
Network Mask: /0
    Metric Type: 1
    TOS: 0
    Metric: 16777215
    Forward Address: 0.0.0.0
    External Route Tag: 0

LS age: 70
Options: (No TOS-Capability, No DC, No Type 7/5 translation)
LS type: NSSA External
Link State ID: 0.0.0.0
Advertising Router: 2.2.2.2
LS Seq Number: 0x80000001
Checksum: 0x2526
Length: 36
Network Mask: /0
    Metric Type: 1
```

```
    TOS: 0
    Metric: 0
    Forward Address: 0.0.0.0
    External Route Tag: 0

LS age: 65
Options: (No TOS-Capability, No DC, No Type 7/5 translation)
LS type: NSSA External
Link State ID: 12.1.1.0
Advertising Router: 2.2.2.2
LS Seq Number: 0x80000001
Checksum: 0xBDEA
Length: 36
Network Mask: /24
    Metric Type: 2
    TOS: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0

LS age: 65
Options: (No TOS-Capability, No DC, No Type 7/5 translation)
LS type: NSSA External
Link State ID: 13.1.1.0
Advertising Router: 2.2.2.2
LS Seq Number: 0x80000001
Checksum: 0xB0F6
Length: 36
Network Mask: /24
    Metric Type: 2
    TOS: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0

LS age: 65
Options: (No TOS-Capability, No DC, No Type 7/5 translation)
LS type: NSSA External
Link State ID: 14.1.1.0
Advertising Router: 2.2.2.2
LS Seq Number: 0x80000001
Checksum: 0xA303
Length: 36
Network Mask: /24
    Metric Type: 2
    TOS: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

Supported Releases 10.2.0E or later

show ip ospf database opaque-area

Displays information about the opaque-area Type 10 LSA.

Syntax `show ip ospf [process-id] [vrf vrf-name] database opaque-area`

Parameters

- *process-id* — (Optional) Displays the opaque-area Type 10 information for an OSPF process ID. If you do not enter a process ID, this command applies only to the first OSPF process.
- `vrf vrf-name` — (Optional) Displays the opaque-area Type 10 information for an OSPF process ID corresponding to a VRF.

Default Not configured

Command Mode EXEC

Usage Information

- `LS Age` — Displays the LS age.
- `Options` — Displays the optional capabilities available on the router.
- `LS Type` — Displays the LS type.
- `Link State ID` — Identifies the router ID.
- `Advertising Router` — Identifies the advertising router's ID.
- `LS Seq Number` — Identifies the LS sequence number. This identifies old or duplicate LSAs.
- `Checksum` — Displays the Fletcher checksum of an LSA's complete contents.
- `Length` — Displays the LSA length in bytes.
- `Opaque Type` — Identifies the Opaque type field, the first 8 bits of the LS ID.
- `Opaque ID` — Identifies the Opaque type-specific ID, the remaining 24 bits of the LS ID.

Example

```
OS10# show ip ospf database opague-area
       OSPF Router with ID (1.1.1.1) (Process ID 100)

          Type-10 Area Local Opaque (Area 0.0.0.1)

LS age: 3600
Options: (No TOS-Capability, No DC)
LS type: Type-10 Area Local Opaque
Link State ID: 8.1.1.2
Advertising Router: 2.2.2.2
LS Seq Number: 0x80000008
Checksum: 0x83B8
Length: 28
Opaque Type: 8
Opaque ID: 65794
!!
!
```

Supported Releases 10.2.0E or later

show ip ospf database opaque-as

Displays information about the opaque-as Type 11 LSAs.

Syntax `show ip ospf [process-id] opaque—as`

Parameters *process-id* — (Optional) Displays opaque-as Type 11 LSA information for a specified OSPF process ID. If you do not enter a process ID, this command applies only to the first OSPF process.

Default Not configured

Command Mode EXEC

Usage Information

- `LS Age` — Displays the LS age.
- `Options` — Displays the optional capabilities available on the router.
- `LS Type` — Displays the LS type.
- `Link State ID` — Identifies the router ID.
- `Advertising Router` — Identifies the advertising router's ID.
- `LS Seq Number` — Identifies the LS sequence number. This identifies old or duplicate LSAs.
- `Checksum` — Displays the Fletcher checksum of an LSA's complete contents.
- `Length` — Displays the LSA length in bytes.
- `Opaque Type` — Identifies the Opaque type field, the first 8 bits of the LS ID.

- `Opaque ID` — Identifies the Opaque type-specific ID, the remaining 24 bits of the LS ID.

Example

```
OS10# show ip ospf 100 database opaque-as

  OSPF Router with ID (1.1.1.1) (Process ID 100)

          Type-11 AS Opaque

LS age: 3600
Options: (No TOS-Capability, No DC)
LS type: Type-11 AS Opaque
Link State ID: 8.1.1.3
Advertising Router: 2.2.2.2
LS Seq Number: 0x8000000D
Checksum: 0x61D3
Length: 36
Opaque Type: 8
Opaque ID: 65795
```

Supported Releases 10.2.0E or later

show ip ospf database opaque-link

Displays information about the opaque-link Type 9 LSA.

Syntax `show ip ospf [process-id] [vrf vrf-name] database opaque-link`

Parameters

- *process-id* — (Optional) Displays the opaque-link Type 9 LSA information for an OSPF process ID. If you do not enter a process ID, this command applies only to the first OSPF process.
- `vrf vrf-name` — (Optional) Displays the opaque-link Type 9 LSA information for an OSPF process ID corresponding to a VRF.

Default Not configured

Command Mode EXEC

Usage Information

- `LS Age` — Displays the LS age.
- `Options` — Displays the optional capabilities available on the router.
- `LS Type` — Displays the LS type.
- `Link State ID` — Identifies the router ID.
- `Advertising Router` — Identifies the advertising router's ID.
- `LS Seq Number` — Identifies the LS sequence number. This identifies old or duplicate LSAs.
- `Checksum` — Displays the Fletcher checksum of an LSA's complete contents.
- `Length` — Displays the LSA length in bytes.
- `Opaque Type` — Identifies the Opaque type field, the first 8 bits of the LS ID.
- `Opaque ID` — Identifies the Opaque type-specific ID, the remaining 24 bits of the LS ID.

Example

```
OS10# show ip ospf 100 database opaque-link
       OSPF Router with ID (1.1.1.1) (Process ID 100)

          Type-9 Link Local Opaque (Area 0.0.0.1)

LS age: 3600
Options: (No TOS-Capability, No DC)
LS type: Type-9 Link Local Opaque
Link State ID: 8.1.1.1
Advertising Router: 2.2.2.2
LS Seq Number: 0x80000007
```

```
Checksum: 0x9DA1
Length: 28
Opaque Type: 8
Opaque ID: 65793
```

Supported Releases 10.2.0E or later

show ip ospf database router

Displays information about the router Type 1 LSA.

Syntax `show ip ospf process-id [vrf vrf-name] database router`

Parameters

- *process-id* — (Optional) Displays the router Type 1 LSA for an OSPF process ID. If you do not enter a process ID, this command applies only to the first OSPF process.
- `vrf` *vrf-name* — (Optional) Displays the router Type 1 LSA for an OSPF process ID corresponding to a VRF.

Default Not configured

Command Mode EXEC

Usage Information Output:

- `LS age`—Displays the LS age.
- `Options`—Displays optional capabilities.
- `LS Type`—Displays the LS type.
- `Link State ID`—Identifies the router ID.
- `Advertising Router`—Identifies the advertising router's ID.
- `LS Seq Number`—Identifies the LS sequence number. This identifies old or duplicate LSAs.
- `Checksum`—Displays the Fletcher checksum of an LSA's complete contents.
- `Length`—Displays the LSA length in bytes.
- `TOS`—Displays the ToS options. The only option available is zero.
- `Metric`—Displays the LSA metric.

Example

```
OS10# show ip ospf 10 database router

       OSPF Router with ID (111.2.1.1) (Process ID 10)

                Router (Area 0.0.0.0)

  LS age: 1419
  Options: (No TOS-capability, No DC, E)
  LS type: Router
  Link State ID: 111.2.1.1
  Advertising Router: 111.2.1.1
  LS Seq Number: 0x8000000d
  Checksum: 0x9bf2
  Length: 60
  AS Boundary Router
   Number of Links: 3

    Link connected to: a Transit Network
     (Link ID) Designated Router address: 110.1.1.2
     (Link Data) Router Interface address: 110.1.1.1
     Number of TOS metric: 0
      TOS 0 Metric: 1

    Link connected to: a Transit Network
     (Link ID) Designated Router address: 111.1.1.1
```

```
     (Link Data) Router Interface address: 111.1.1.1
     Number of TOS metric: 0
      TOS 0 Metric: 1

    Link connected to: a Transit Network
     (Link ID) Designated Router address: 111.2.1.1
     (Link Data) Router Interface address: 111.2.1.1
     Number of TOS metric: 0
      TOS 0 Metric: 1
```

Supported Releases 10.2.0E or later

show ip ospf database summary

Displays the network summary Type 3 LSA routing information.

Syntax `show ip ospf [process-id] [vrf vrf-name] database summary`

Parameters

- *process-id*—(Optional) Displays LSA information for a specific OSPF process ID. If you do not enter a process ID, this command applies only to the first OSPF process.
- `vrf` *vrf-name* — (Optional) Displays LSA information for a specified OSPF process ID corresponding to a VRF.

Default Not configured

Command Mode EXEC

Usage Information

- `LS Age`—Displays the LS age.
- `Options`—Displays the optional capabilities available on the router.
- `LS Type`—Displays the LS type.
- `Link State ID`—Identifies the router ID.
- `Advertising Router`—Identifies the advertising router's ID.
- `LS Seq Number`—Identifies the LS sequence number. This identifies old or duplicate LSAs.
- `Checksum`—Displays the Fletcher checksum of an LSA's complete contents.
- `Length`—Displays the LSA length in bytes.
- `Network Mask`—Identifies the network mask implemented on the area.
- `TOS`—Displays the ToS options. The only option available is zero.
- `Metric`—Displays the LSA metric.

Example

```
OS10# show ip ospf 10 database summary
     OSPF Router with ID (111.2.1.1) (Process ID 10)

                Summary Network (Area 0.0.0.0)

  LS age: 623
  Options: (No TOS-capability, No DC)
  C: Summary Network
  Link State ID: 115.1.1.0
  Advertising Router: 111.111.111.1
  LS Seq Number: 0x800001e8
  Checksum: 0x4a67
  Length: 28
  Network Mask: /24
      TOS: 0 Metric: 0
```

Supported Releases 10.2.0E or later

show ip ospf interface

Displays the configured OSPF interfaces. You must enable OSPF to display output.

Syntax

`show ip ospf interface [process-id] [vrf vrf-name] interface` or `show ip ospf [process-id] [vrf vrf-name] interface [interface]`

Parameters

- *process-id* — (Optional) Displays information for an OSPF process ID. If you do not enter a process ID, this command applies only to the first OSPF process.
- `vrf` *vrf-name* — (Optional) Displays information for an OSPF instance corresponding to a VRF.
- *interface* — (Optional) Enter the interface information:
  - `ethernet` — Enter the Ethernet interface information, from 1 to 48.
  - `port channel` — Enter the port-channel interface number, from 1 to 128.
  - `vlan` — Enter the VLAN interface number, from 1 to 4093.

Default Not configured

Command Mode EXEC

Example

```
OS10# show ip ospf 10 interface
ethernet1/1/1 is up, line protocol is up
  Internet Address 110.1.1.1/24, Area 0.0.0.0
  Process ID 10, Router ID 1.1.1.1, Network Type broadcast, Cost: 10
  Transmit Delay is 1 sec, State WAIT, Priority 1
  BFD enabled(Interface level) Interval 300 Min_rx 300 Multiplier 3 Role Active
  Designated Router (ID) , Interface address 0.0.0.0
  Backup Designated router (ID) , Interface address 0.0.0.0
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Neighbor Count is 0, Adjacent neighbor count is 0
```

Supported Releases 10.2.0E or later

show ip ospf routes

Displays OSPF routes received from neighbors along with parameters such as cost, next-hop, area, interface, and type of route.

Syntax

`show ip ospf [process-id] [vrf vrf-name] routes [prefix IP-prefix]`

Parameters

- *process-id* — (Optional) Enter OSPFv2 process ID to view information specific to the ID.
- `vrf` *vrf-name* — (Optional) Enter the keyword vrf followed by the name of the VRF to display the routes calculated by OSPF in the configured VRF.
- *IP-prefix* — (Optional) Specify an IP address to view information specific to the IP address.

Default None

Command Mode EXEC

Usage Information Displays the cost metric for each neighbor and interfaces.

Example

```
OS10# show ip ospf 10 routes
Prefix      Cost  Nexthop   Interface   Area     Type
110.1.1.0   1     0.0.0.0   vlan3050   0.0.0.0  intra-area
111.1.1.0   1     0.0.0.0   vlan3051   0.0.0.0  intra-area
111.2.1.0   1     0.0.0.0   vlan3053   0.0.0.0  intra-area
```

Supported Releases 10.2.0E or later

show ip ospf statistics

Displays OSPF traffic statistics.

Syntax

- `show ip ospf [instance-number] [vrf vrf-name] statistics [interface interface]`

Parameters

- *instance-number* — (Optional) Enter an OSPF instance number, from 1 to 65535.
- `vrf` *vrf-name* — (Optional) Enter the keyword `vrf` followed by the name of the VRF to display OSPF traffic statistics corresponding to that VRF.
- `interface` *interface* — (Optional) Enter the interface information:
  - `ethernet` *node/slot/port[:subport]* — Enter an Ethernet port interface.
  - `port-channel` *number* — Enter the port-channel interface number, from 1 to 128.
  - `vlan` *vlan-id* — Enter the VLAN ID number, from 1 to 4093.

Default Not configured

Command Mode EXEC

Usage Information This command displays OSPFv2 traffic statistics for a specified instance or interface, or for all OSPFv2 instances and interfaces.

Example

```
OS10# show ip ospf 10 statistics
Interface vlan3050
  Receive Statistics
     rx-invalid          0    rx-invalid-bytes     0
     rx-hello            0    rx-hello-bytes       0
     rx-db-des           0    rx-db-des-bytes      0
     rx-ls-req           0    rx-ls-req-bytes      0
     rx-ls-upd           0    rx-ls-upd-bytes      0
     rx-ls-ack           0    rx-ls-ack-bytes      0
  Transmit Statistics
     tx-failed           0    tx-failed-bytes      0
     tx-hello            0    tx-hello-bytes       0
     tx-db-des           0    tx-db-des-bytes      0
     tx-ls-req           0    tx-ls-req-bytes      0
     tx-ls-upd           0    tx-ls-upd-bytes      0
     tx-ls-ack           0    tx-ls-ack-bytes      0
  Error packets (Receive statistics)
     bad-src             0   dupe-id               0   hello-err           0
     mtu-mismatch        0   nbr-ignored           0   wrong-proto         0
     resource-err        0   bad-lsa-len           0   lsa-bad-type        0
     lsa-bad-len         0   lsa-bad-cksum         0   auth-fail           0
     netmask-mismatch    0   hello-tmr-mismatch    0   dead-ivl-mismatch   0
     options-mismatch    0   nbr-admin-down        0   own-hello-drop      0
     self-orig             0   wrong-length        0   checksum-error      0
     version-mismatch      0   area-mismatch       0
```

Supported Releases 10.2.0E or later

show ip ospf topology

Displays routers that directly connect to OSPF areas.

Syntax `show ip ospf [process-id] [vrf vrf-name] topology`

Parameters

- *process-id* — (Optional) Displays OSPF process information. If you do not enter a process ID, this applies only to the first OSPF process.

- `vrf vrf-name` — (Optional) Displays the routers in the directly connected OSPF areas in the configured VRF.

Default Not configured

Command Mode EXEC

Usage Information The "E" flag output indicates the router listed is an ASBR. The "B" flag indicates that the router listed is an ABR. If the Flag field shows both E and B, it indicates that the listed router is both an ASBR and an ABR.

Example

```
OS10# show ip ospf 10 topology

 Router ID      Flags   Cost  Nexthop     Interface  Area
111.111.111.1   -/B/-/   1    111.1.1.2   Vl 3051    0
111.111.111.2   -/B/-/   1    111.2.1.2   Vl 3053    0
112.2.1.1       E/-/-/   1    110.1.1.2   Vl 3050    0
112.112.112.1   -/B/-/   2    110.1.1.2   Vl 3050    0
112.112.112.2   -/B/-/   2    110.1.1.2   Vl 3050    0
```

Supported Releases 10.2.0E or later

summary-address

Configures a summary address for an ASBR to advertise one external route as an aggregate for all redistributed routes covered by a specified address range.

Syntax `summary-address ip-address/mask [not-advertise | tag tag-value]`

Parameters

- `ip-address/mask`—Enter the IP address to summarize along with the mask.
- `not-advertise`—(Optional) Suppresses IP addresses that do not match the network prefix/mask.
- `tag-value`—(Optional) Enter a value to match the routes redistributed through a route map, from 1 to 65535.

Default) Not configured

Command Mode ROUTER-OSPF

Usage Information The `no` version of this command disables the summary address.

Example

```
OS10(config)# router ospf 100
OS10(config-router-ospf-100)# summary-address 10.0.0.0/8 not-advertise
```

Supported Releases 10.3.0E or later

timers lsa arrival

Configures the LSA acceptance intervals.

Syntax `timers lsa arrival arrival-time`

Parameters `arrival-time` — Set the interval between receiving the LSA in milliseconds, from 0 to 600,000.

Default 1000 milliseconds

Command Mode ROUTER-OSPF

Usage Information Setting the LSA arrival time between receiving the LSA repeatedly ensures that the system gets enough time to accept the LSA. The `no` version of this command resets the value to the default.

Example

```
OS10(config)# router ospf 10
OS10(conf-router-ospf-10)# timers lsa arrival 2000
```

Supported Releases 10.2.0E or later

timers spf

Enables shortest path first (SPF) throttling to delay an SPF calculation when a topology change occurs.

Syntax `timers spf [start-time [hold-time [max-wait]]]`

Parameters

- *start-time* — Sets the initial SPF delay in milliseconds, from 1 to 600000; default 1000.
- *hold-time* — Sets the additional hold time between two SPF calculations in milliseconds, from 1 to 600000; default 10000.
- *max-wait* — Sets the maximum wait time between two SPF calculations in milliseconds, from 1 to 600000; default 10000.

Default

- *start-time* — 1000 milliseconds
- *hold-time* — 10000 milliseconds
- *max-wait* — 10000 milliseconds

Command Mode ROUTER-OSPF

Usage Information By default, SPF timers are disabled in an OSPF instance.

Use SPF throttling to delay SPF calculations during periods of network instability. In an OSPF network, a topology change event triggers an SPF calculation after a start time. When the start timer finishes, a hold time may delay the next SPF calculation for an additional time. When the hold timer is running:

- Each time a topology change occurs, the SPF calculation delays for double the configured hold time up to maximum wait time.
- If no topology change occurs, an SPF calculation performs and the hold timer is reset to its configured value.

If you do not specify a start-time, hold-time, or max-wait value, the default values are used. The `no` version of this command removes the configured SPF timers and disables SPF throttling in an OSPF instance.

Example

```
OS10(config)# router ospf 100
OS10(config-router-ospf-100)# timers spf 1200 2300 3400
OS10(config-router-ospf-100)# do show ip ospf

Routing Process ospf 100 with ID 12.1.1.1
Supports only single TOS (TOS0) routes
It is Flooding according to RFC 2328
SPF schedule delay 1200 msecs, Hold time between two SPFs 2300 msecs
Convergence Level 0
Min LSA origination 0 msec, Min LSA arrival 1000 msec
Min LSA hold time 5000 msec, Max LSA wait time 5000 msec
Number of area in this router is 1, normal 1 stub 0 nssa 0
Area (0.0.0.1)
Number of interface in this area is 1
SPF algorithm executed 1 times
```

Supported Releases 10.4.0E(R1) or later

timers throttle lsa all

Configures the LSA transmit intervals.

Syntax `timers lsa all [start-interval | hold-interval | max-interval]`

Parameters

- *start-interval* — Sets the minimum interval between initial sending and re-sending the same LSA in milliseconds, from 0 to 600,000.
- *hold-interval* — Sets the next interval to send the same LSA in milliseconds. This is the time between sending the same LSA after the start-interval is attempted, from 1 to 600,000.
- *max-interval* — Sets the maximum amount of time the system waits before sending the LSA in milliseconds, from 1 to 600,000.

.

Default

- start-interval — 0 milliseconds
- hold-interval — 5000 milliseconds
- max-interval — 5000 milliseconds

Command Mode ROUTER-OSPF

Usage Information The `no` version of this command removes the LSA transmit timer.

Example

```
OS10(config)# router ospf 10
OS10(conf-router-ospf-10)# timers throttle lsa all 100 300 1000
```

Supported Releases 10.2.0E or later

OSPFv3

OSPFv3 is an IPv6 link-state routing protocol that supports IPv6 unicast address families (AFs). OSPFv3 is disabled by default. You must configure at least one interface, either physical or Loopback. The OSPF process automatically starts when OSPFv3 is enabled for one or more interfaces. Any area besides *area 0* can have any number ID assigned to it.

Enable OSPFv3

1 Enable OSPFv3 globally and configure an OSPFv3 instance in CONFIGURATION mode.

```
router ospfv3 instance-number
```

2 Enter the interface information to configure the interface for OSPFv3 in INTERFACE mode.

```
interface ethernet node/slot/port[:subport]
```

3 Enable the interface in INTERFACE mode.

```
no shutdown
```

4 Disable the default switchport configuration and remove it from an interface or a LAG port in INTERFACE mode.

```
no switchport
```

5 Enable the OSPFv3 on an interface in INTERFACE mode.

```
ipv6 ospfv3 process-id area area-id
```

- *process-id* — Enter the OSPFv3 process ID for a specific OSPFv3 process, from 1 to 65535.
- *area-id* — Enter the OSPF area ID as an IP address in A.B.C.D format or number, from 1 to 65535.

Enable OSPFv3

```
OS10(config)# router ospfv3 100
OS10(config-router-ospfv3-100)# exit
OS10(config)#  interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# ipv6 ospfv3 300 area 0.0.0.0
```

Enable OSPFv3 in a non-default VRF instance

1 Create the non-default VRF instance in which you want to enable OSPFv3:

```
ip vrf vrf-name
```

CONFIGURATION Mode

2 Enable OSPFv3 in the non-default VRF instance that you created earlier and configure an OSPFv3 instance in VRF CONFIGURATION mode.

```
router ospfv3 instance-number vrf vrf-name
```

3 Enter the interface information to configure the interface for OSPFv3 in INTERFACE mode.

```
interface ethernet node/slot/port[:subport]
```

4 Enable the interface in INTERFACE mode.

```
no shutdown
```

5 Disable the default switchport configuration and remove it from an interface or a LAG port in INTERFACE mode.

```
no switchport
```

6 Associate the interface with the non-default VRF instance that you created earlier.

```
ip vrf forwarding vrf-name
```

7 Enable the OSPFv3 on an interface.

```
ipv6 ospfv3 process-id area area-id
```

- *process-id* — Enter the OSPFv3 process ID for a specific OSPFv3 process, from 1 to 65535.
- *area-id* — Enter the OSPF area ID as an IP address in A.B.C.D format or number, from 1 to 65535.

Enable OSPFv3

```
OS10(config)# ip vrf vrf-blue
OS10(config-vrf-blue)# router ospfv3 100 vrf vrf-blue
OS10(config-router-ospfv3-100)# exit
OS10(config)#  interface ethernet 1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/2)# ip vrf forwarding vrf-blue
OS10(conf-if-eth1/1/1)# ipv6 ospfv3 300 area 0.0.0.0
```

Assign Router ID

You can assign a router ID for the OSPFv3 process. Configure an arbitrary value in the IP address format for each router. Each router ID must be unique. Use the fixed router ID for the active OSPFv3 router process. Changing the router ID brings down the existing OSPFv3 adjacency. The new router ID becomes effective immediately.

- Assign the router ID for the OSPFv3 process in ROUTER-OSPFv3 mode.

```
router-id ip-address
```

Assign router ID

```
OS10(config)# router ospfv3 100
OS10(config-router-ospfv3-100)# router-id 10.10.1.5
```

View OSPFv3 Status

```
OS10# show ipv6 ospf
Routing Process ospfv3 100 with ID 10.10.1.5
It is an Area Border Router
Min LSA origination 5000 msec, Min LSA arrival 1000 msec
Min LSA hold time 0 msec, Max LSA wait time 0 msec
Number of area in this router is 2, normal 2 stub 0 nssa
   Area  (0.0.0.0)
       Number of interface in this area is 1
       SPF algorithm executed 42 times
   Area  (0.0.0.1)
       Number of interface in this area is 1
       SPF algorithm executed 42 times
```

Configure Stub Areas

Type 5 LSAs are not flooded into stub areas. The ABR advertises a default route into the stub area where it is attached. Stub area routers use the default route to reach external destinations.

1 Enable OSPFv3 routing and enter ROUTER-OSPFv3 mode, from 1 to 65535.

```
router ospfv3 instance number
```

2 Configure an area as a stub area in ROUTER-OSPFv3 mode.

```
area area-id stub [no-summary]
```

- *area-id* — Enter the OSPFv3 area ID as an IP address in A.B.C.D format or number, from 1 to 65535.
- no-summary — (Optional) Enter to prevent an ABR from sending summary LSAs into the stub area.

Configure Stub Area

```
OS10(config)# router ospfv3 10
OS10(conf-router-ospf-10)# area 10.10.5.1 stub no-summary
```

View Stub Area Configuration

```
OS10# show running-configuration ospfv3
!
interface ethernet1/1/3
ipv6 ospf 65 area 0.0.0.2
!
router ospfv3 65
area 0.0.0.2 stub no-summary


OS10# show ipv6 ospf database
       OSPF Router with ID (199.205.134.103) (Process ID 65)

Router Link States (Area 0.0.0.2)

ADV Router      Age       Seq#        Fragment ID Link count Bits
-------------------------------------------------------------------
199.205.134.103 32        0x80000002  0           1
202.254.156.15  33        0x80000002  0           1          B

Net Link States (Area 0.0.0.2)

ADV Router      Age       Seq#        Link ID   Rtr count
-----------------------------------------------------------
202.254.156.15  38        0x80000001  12        2

Inter Area Prefix Link States (Area 0.0.0.2)

ADV Router      Age       Seq#        Prefix
-----------------------------------------------------------------
202.254.156.15  93        0x80000001  ::/0
```

```
Intra Area Prefix Link States  (Area 0.0.0.2)

ADV Router      Age       Seq#        Link ID   Ref-lstype Ref-LSID
-------------------------------------------------------------------
202.254.156.15  34        0x80000003  65536     0x2002     12

Link (Type-8) Link States (Area 0.0.0.2)

ADV Router      Age       Seq#        Link ID   Interface
-----------------------------------------------------------------
199.205.134.103 42        0x80000001  12        ethernet1/1/3
202.254.156.15  54        0x80000001  12        ethernet1/1/3
```

Enable Passive Interfaces

A passive interface is one that does not send or receive routing information. Configuring an interface as a passive interface suppresses both the receiving and sending routing updates.

Although the passive interface does not send or receive routing updates, the network on that interface is included in OSPF updates sent through other interfaces. You can remove an interface from passive interfaces using the `no ipv6 ospf passive` command.

1 Enter an interface type in INTERFACE mode.

```
interface ethernet node/slot/port[:subport]
```

2 Configure the interface as a passive interface in INTERFACE mode.

```
ipv6 ospf passive
```

Configure Passive Interfaces

```
OS10(config)# interface ethernet 1/1/6
OS10(conf-if-eth1/1/6)# ipv6 ospf passive
```

View Passive Interfaces

```
OS10# show running-configuraiton
!!!
!!
interface ethernet1/1/1
 ip address 10.10.10.1/24
 no switchport
 no shutdown
 ipv6 ospf 100 area 0
 ipv6 ospf passive
!!
!
```

Interface OSPFv3 Parameters

Interface parameter values must be consistent across all interfaces to avoid routing errors. For example, set the same time interval for the hello packets on all routers in the OSPF network to prevent misconfiguration of OSPF neighbors.

1 Enter the interface to change the OSPFv3 parameters in CONFIGURATION mode.

```
interface interface-name
```

2 Change the cost associated with OSPFv3 traffic on the interface in INTERFACE mode, from 1 to 65535, default depends on the interface speed.

```
ipv6 ospf cost
```

3 Change the time interval the router waits before declaring a neighbor dead in INTERFACE mode, from 1 to 65535, default 40. The dead interval must be four times the hello interval. The dead interval must be the same on all routers in the OSPFv3 network.

```
ipv6 ospf dead-interval seconds
```

4 Change the time interval in seconds between hello-packet transmission in INTERFACE mode, from 1 to 65535, default 10. The hello interval must be the same on all routers in the OSPFv3 network.

```
ipv6 ospf hello-interval seconds
```

5 Change the priority of the interface, which determines the DR for the OSPFv3 broadcast network in INTERFACE mode, from 0 to 255, default 1.

```
ipv6 ospf priority number
```

Change OSPFv3 Interface Parameters

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ipv6 ospf hello-interval 5
OS10(conf-if-eth1/1/1)# ipv6 ospf dead-interval 20
OS10(conf-if-eth1/1/1)# ipv6 ospf priority 4
```

View OSPFv3 Interface Parameters

```
OS10# show ipv6 ospf interface
ethernet1/1/1 is up, line protocol is up
  Link Local Address fe80::20c:29ff:fe0a:d59/64, Interface ID 5
  Area 0.0.0.0, Process ID 200, Instance ID 0, Router ID 10.0.0.2
  Network Type broadcast, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router on this network is 2.2.2.2
  Backup Designated router on this network is 10.0.0.2 (local)
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 2.2.2.2(Designated Router)
```

Default route

You can generate an external default route and distribute the default information to the OSPFv3 routing domain.

- Generate the default route, using the `default-information originate [always]` command in ROUTER-OSPFv3 mode.

Configure default route

```
OS10(config)# router ospfv3 100
OS10(config-router-ospf-100)# default-information originate always
```

View default route configuration

```
OS10(config-router-ospf-100)# show configuration
!
router ospfv3 100
 default-information originate always
```

OSPFv3 IPsec authentication and encryption

Unlike OSPFv2, OSPFv3 does not have authentication fields in its protocol header to provide security. To provide authentication and confidentiality, OSPFv3 uses IP Security (IPsec) — a collection of security protocols for authenticating and encrypting data packets. OS10 OSPFv3 supports IPsec using the IPv6 authentication header (AH) or IPv6 encapsulating security payload (ESP).

- AH authentication verifies that data is not altered during transmission and ensures that users are communicating with the intended individual or organization. The authentication header is inserted after the IP header with a value of 51. MD5 and SHA1 authentication types are supported; encrypted and unencrypted keys are supported.
- ESP encryption encapsulates data, enabling data protection that follows in the datagram. The ESP extension header is inserted after the IP header and before the next layer protocol header. 3DES, DES, AES-CBC, and NULL encryption algorithms are supported; encrypted and unencrypted keys are supported.

Apply IPsec authentication or encryption on a physical, port-channel, or VLAN interface or in an OSPFv3 area. Each configuration consists of a security policy index (SPI) and the OSPFv3 packets validation key. After you configure an IPsec protocol for OSPFv3, IPsec operation is invisible to the user.

You can only enable one authentication or encryption security protocol at a time on an interface or for an area. Enable IPsec AH using the `ipv6 ospf authentication` command; enable IPsec ESP with the `ipv6 ospf encryption` command.

- A security policy configured for an area is inherited on all interfaces in the area by default.
- A security policy configured on an interface overrides any area-level configured security for the area where the interface is assigned.
- The configured authentication or encryption policy applies to all OSPFv3 packets transmitted on the interface or in the area. The IPsec security associations are the same on inbound and outbound traffic on an OSPFv3 interface.
- There is no maximum AH or ESP header length because the headers have fields with variable lengths.

Configure IPsec authentication on interfaces

Prerequisite: Before you enable IPsec authentication on an OSPFv3 interface, first enable IPv6 unicast routing globally, then enable OSPFv3 on the interface, and assign it to an area.

The SPI value must be unique to one IPsec authentication or encryption security policy on the router. You cannot configure the same SPI value on another interface even if it uses the same authentication or encryption algorithm.

You cannot use an IPsec MD5 or SHA-1 authentication type and the `null` setting at same time on an interface. These settings are mutually exclusive.

- Enable IPsec authentication for OSPFv3 packets in Interface mode.

  ```
  ipv6 ospf authentication {null | ipsec spi number {MD5 | SHA1} key}
  ```

  - `null` — Prevent an authentication policy configured for the area to be inherited on the interface. Only use this parameter if you configure IPsec area authentication.
  - `ipsec spi number` — Enter a unique security policy index (SPI) value, from 256 to 4294967295.
  - `md5` — Enable message digest 5 (MD5) authentication.
  - `sha1` — Enable secure hash algorithm 1 (SHA-1) authentication.
  - *`key`* — Enter the text string used in the authentication type. All neighboring OSPFv3 routers must share the key to exchange information. Only a non-encrypted key is supported. For MD5 authentication, the non-encrypted key must be 32 plain hex digits. For SHA-1 authentication, the non-encrypted key must be 40 hex digits. An encrypted key is not supported.

To delete an IPsec authentication policy, use the `no ipv6 ospf authentication ipsec spi number` or `no ipv6 ospf authentication null` command.

Configure IPsec authentication on interface

```
OS10(conf-if-eth1/1/1)# ipv6 ospf authentication ipsec spi 400 md5
12345678123456781234567812345678
OS10(conf-if-eth1/1/1)# show configuration
!
interface ethernet1/1/1
ipv6 ospf authentication ipsec spi 400 md5 12345678123456781234567812345678
no switchport
no shutdown
ipv6 address 1::1/64
```

IPsec encryption on interfaces

Prerequisite: Before you enable IPsec encryption on an OSPFv3 interface, enable IPv6 unicast routing globally, enable OSPFv3 on the interface, and assign it to an area.

When you configure encryption on an interface, both IPsec encryption and authentication are enabled. You cannot configure encryption if you have already configured an interface for IPsec authentication using the `ipv6 ospf authentication ipsec` command. To configure encryption, you must first delete the authentication policy.

- Enable IPsec encryption for OSPFv3 packets in Interface mode.

  ```
  ipv6 ospf encryption ipsec spi number esp encryption-type
  key authentication-type key
  ```

- `ipsec spi number` — Enter a unique security policy index (SPI) value, from 256 to 4294967295.
- `esp encryption-type key` — Enter the encryption algorithm used with ESP (3DES, DES, AES-CBC, or NULL). For AES-CBC, only the AES-128 and AES-192 ciphers are supported.
- *key* — Enter the text string used in the encryption algorithm. All neighboring OSPFv3 routers must share the key to decrypt information. Only a non-encrypted key is supported. Required lengths of the non-encrypted key are: 3DES — 48 hex digits; DES — 16 hex digits; AES-CBC — 32 hex digits for AES-128 and 48 hex digits for AES-192.
- *authentication-type key* — Enter the encryption authentication MD5 or SHA1 algorithm to use.
- *key* — Enter the text string used in the authentication algorithm. All neighboring OSPFv3 routers must share the key to exchange information. Only a non-encrypted key is supported. For MD5 authentication, the non-encrypted key must be 32 plain hex digits. For SHA1 authentication, the non-encrypted key must be 40 hex digits. An encrypted key is not supported.

To delete an IPsec encryption policy, use the `no ipv6 ospf encryption ipsec spi number` or `no ipv6 ospf encryption null` command.

Configure IPsec encryption on interface

```
OS10(conf-if-eth1/1/1)# ipv6 ospf encryption ipsec spi 500 esp des 1234567812345678 md5
12345678123456781234567812345678
OS10(conf-if-eth1/1/1)# show configuration
!
interface ethernet1/1/1
ipv6 ospf encryption ipsec spi 500 esp des 1234567812345678 md5 12345678123456781234567812345678
no switchport
no shutdown
ipv6 address 1::1/64
```

Configure IPsec authentication for OSPFv3 area

Prerequisite: Before you enable IPsec authentication for an OSPFv3 area, enable OSPFv3 globally on the router.

- Enable IPsec authentication for OSPFv3 packets in an area in Router-OSPFv3 mode.

  `area area-id  authentication ipsec spi number {MD5 | SHA1} key`

  - `area area-id` — Enter an area ID as a number or IPv6 prefix.
  - `ipsec spi number` — Enter a unique security policy index (SPI) value, from 256 to 4294967295.
  - `md5` — Enable message digest 5 (MD5) authentication.
  - `sha1` — Enable secure hash algorithm 1 (SHA1) authentication.
  - *key* — Enter the text string used in the authentication type. All OSPFv3 routers in the area share the key to exchange information. Only a non-encrypted key is supported. For MD5 authentication, the non-encrypted key must be 32 plain hex digits. For SHA1 authentication, the non-encrypted key must be 40 hex digits. An encrypted key is not supported.

To delete an IPsec area authentication policy, use the `no area area-id authentication ipsec spi number` command.

Configure IPsec authentication for an OSPFv3 area

```
OS10(config-router-ospfv3-100)# area 1 authentication ipsec spi 400 md5
12345678123456781234567812345678
OS10(config-router-ospfv3-100)# show configuration
!
router ospfv3 100
area 0.0.0.1 authentication ipsec spi 400 md5 12345678123456781234567812345678
```

IPsec encryption for OSPV3 area

Prerequisite: Before you enable IPsec encryption for an OSPFv3 area, first enable OSPFv3 globally on the router.

When you configure encryption at the area level, both IPsec encryption and authentication are enabled. You cannot configure encryption if you have already configured an IPsec area authentication using the `area ospf authentication ipsec` command. To configure encryption, you must first delete the authentication policy.

- Enable IPsec encryption for OSPFv3 packets in an area in Router-OSPFv3 mode.

  ```
  area area-id  encryption ipsec spi number esp encryption-type key
  authentication-type key
  ```

  - `area area-id` — Enter an area ID as a number or IPv6 prefix.
  - `ipsec spi number` — Enter a unique security policy index (SPI) value, from 256 to 4294967295.
  - `esp encryption-type` — Enter the encryption algorithm used with ESP (3DES, DES, AES-CBC, or NULL). For AES-CBC, only the AES-128 and AES-192 ciphers are supported.
  - `key` — Enter the text string used in the encryption algorithm. All neighboring OSPFv3 routers must share the key to decrypt information. Only a non-encrypted key is supported. Required lengths of the non-encrypted key are: 3DES — 48 hex digits; DES — 16 hex digits; AES-CBC — 32 hex digits for AES-128 and 48 hex digits for AES-192.
  - `authentication-type` — Enter the encryption authentication MD5 or SHA1 algorithm to use.
  - `key` — Enter the text string used in the authentication algorithm. All neighboring OSPFv3 routers must share the key to exchange information. Only a non-encrypted key is supported. For MD5 authentication, the non-encrypted key must be 32 plain hex digits. For SHA1 authentication, the non-encrypted key must be 40 hex digits. An encrypted key is not supported.

To delete an IPsec encryption policy, use the `no area area-id encryption ipsec spi number` command.

Configure IPsec encryption for OSPFv3 area

```
OS10(config-router-ospfv3-100)# area 1 encryption ipsec spi 401 esp des 1234567812345678 md5
12345678123456781234567812345678
OS10(config-router-ospfv3-100)# show configuration
!
router ospfv3 100
area 0.0.0.1 encryption ipsec spi 401 esp des 1234567812345678 md5
12345678123456781234567812345678
```

Troubleshoot OSPFv3

You can troubleshoot OSPFv3 operations and check questions for typical issues that interrupt a process.

- Is OSPFv3 enabled globally?
- Is OSPFv3 enabled on the interface?
- Are adjacencies established correctly?
- Are the interfaces configured for L3 correctly?
- Is the router in the correct area type?
- Are the OSPF routes included in the OSPF database?
- Are the OSPF routes included in the routing table in addition to the OSPF database?
- Are you able to ping the link-local IPv6 address of adjacent router interface?

Troubleshooting OSPFv3 with show Commands

- View a summary of all OSPF process IDs enabled in EXEC mode.

  ```
  show running-configuration ospfv3
  ```
- View summary information of IP routes in EXEC mode.

  ```
  show ipv6 route summary
  ```
- View summary information for the OSPF database in EXEC mode.

  ```
  show ipv6 ospf database
  ```
- View the configuration of OSPF neighbors connected to the local router in EXEC mode.

  ```
  show ipv6 ospf neighbor
  ```

View OSPF Configuration

```
OS10# show running-configuration ospfv3
!
interface ethernet1/1/1
ip ospf 100 area 0.0.0.0
!
router ospf 100
log-adjacency-changes
```

OSPFv3 Commands

area authentication

Configures authentication for an OSPFv3 area.

Syntax `area area-id authentication ipsec spi number {MD5 | SHA1} key`

Parameters

- `area area-id` — Enter an area ID as a number or IPv6 prefix.
- `ipsec spi number` — Enter a unique security policy index (SPI) value, from 256 to 4294967295.
- `md5` — Enable MD5 authentication.
- `sha1` — Enable SHA1 authentication.
- `key` — Enter the text string used in the authentication type.

Default OSPFv3 area authentication is not configured.

Command Mode ROUTER-OSPFv3

Usage Information

- Before you enable IPsec authentication for an OSPFv3 area, you must enable OSPFv3 globally on each router.
- All OSPFv3 routers in the area must share the same authentication key to exchange information. Only a non-encrypted key is supported. For MD5 authentication, the non-encrypted key must be 32 plain hex digits. For SHA1 authentication, the non-encrypted key must be 40 hex digits. An encrypted key is not supported.

Example

```
OS10(config-router-ospfv3-100)# area 1 authentication ipsec spi 400 md5
12345678123456781234567812345678
```

Supported Releases 10.4.0E(R1) or later

area encryption

Configures encryption for an OSPFv3 area.

Syntax `area area-id encryption ipsec spi number esp encryption-type key authentication-type key`

Parameters

- `area area-id` — Enter an area ID as a number or IPv6 prefix.
- `ipsec spi number` — Enter a unique security policy index number, from 256 to 4294967295.
- `esp encryption-type` — Enter the encryption algorithm used with ESP (3DES, DES, AES-CBC, or NULL). For AES-CBC, only the AES-128 and AES-192 ciphers are supported.
- `key` — Enter the text string used in the encryption algorithm.
- `authentication-type` — Enter the encryption authentication MD5 or SHA1 algorithm to use.
- `key` — Enter the text string used in the authentication algorithm.

Default OSPFv3 area encryption is not configured.

Command Mode ROUTER-OSPFv3

Usage Information

- Before you enable IPsec encryption for an OSPFv3 area, you must enable OSPFv3 globally on each router.
- When you configure encryption at the area level, both IPsec encryption and authentication are enabled. You cannot configure encryption if you have already configured an IPsec area authentication using the `area ospf authentication ipsec` command. To configure encryption, you must first delete the authentication policy.
- All OSPFv3 routers in the area must share the same encryption key to decrypt information. Only a non-encrypted key is supported. Required lengths of the non-encrypted key are: 3DES — 48 hex digits; DES — 16 hex digits; AES-CBC — 32 hex digits for AES-128 and 48 hex digits for AES-192.
- All OSPFv3 routers in the area must share the same authentication key to exchange information. Only a non-encrypted key is supported. For MD5 authentication, the non-encrypted key must be 32 plain hex digits. For SHA1 authentication, the non-encrypted key must be 40 hex digits. An encrypted key is not supported.

Example

```
OS10(config-router-ospfv3-100)# area 1 encryption ipsec spi 401 esp des
1234567812345678 md5
12345678123456781234567812345678
```

Supported Releases 10.4.0E(R1) or later

area stub

Defines an area as the OSPF stub area.

Syntax `area area-id stub [no-summary]`

Parameters

- *area-id*—Set the OSPFv3 area ID as an IP address in A.B.C.D format or number, from 1 to 65535.
- `no-summary`—(Optional) Prevents an ABR from sending summary LAs into the stub area.

Default Not configured

Command Mode ROUTER-OSPFv3

Usage Information The `no` version of this command deletes a stub area.

Example

```
OS10(config)# router ospfv3 10
OS10(conf-router-ospfv3-10)# area 10.10.1.5 stub
```

Supported Releases 10.3.0E or later

auto-cost reference-bandwidth

Calculates default metrics for the interface based on the configured auto-cost reference bandwidth value.

Syntax `auto-cost reference-bandwidth value`

Parameters *value* — Enter the reference bandwidth value to calculate the OSPFv3 interface cost in megabits per second, from 1 to 4294967.

Default 100000

Command Mode ROUTER-OSPFv3

Usage Information The value set by the `ipv6 ospf cost` command in INTERFACE mode overrides the cost resulting from the `auto-cost` command. The `no` version of this command resets the value to the default.

Example

```
OS10(config)# router ospfv3 100
OS10(config-router-ospfv3-100)# auto-cost reference-bandwidth 150
```

Supported Releases 10.3.0E or later

clear ipv6 ospf process

Clears all OSPFv3 routing tables.

Syntax `clear ipv6 ospf {instance-number} [vrf vrf-name] process`

Parameters

- *instance-number* — Enter an OSPFv3 instance number, from 1 to 65535.
- `vrf` *vrf-name* — (Optional) Enter the keyword `vrf` followed by the name of the VRF to clear OSPFv3 processes in that VRF.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# clear ipv6 ospf 3 process
```

Supported Releases 10.3.0E or later

clear ipv6 ospf statistics

Clears OSPFv3 traffic statistics.

Syntax `clear ipv6 ospf [instance-number] [vrf vrf-name] statistics`

Parameters

- *instance-number* — (Optional) Enter an OSPFv3 instance number, from 1 to 65535.
- `vrf` *vrf-name* — (Optional) Enter the keyword vrf followed by the name of the VRF to clear OSPFv3 statistics in that VRF.

Default Not configured

Command Mode EXEC

Usage Information This command clears the OSPFv3 traffic statistics in a specified instance or in all the configured OSPFv3 instances, and resets them to zero.

Example

```
OS10# clear ipv6 ospf 100 statistics
```

Supported Releases 10.4.0E(R1) or later

default-information originate

Generates and distributes a default external route information to the OSPFv3 routing domain.

Syntax `default-information originate [always]`

Parameters `always` — (Optional) Always advertise the default route.

Defaults Disabled

Command Mode ROUTER-OSPFv3

Usage Information The `no` version of this command disables the distribution of default route.

Example

```
OS10(config)# router ospfv3 100
OS10(config-router-ospfv3-100)# default-information originate always
```

Supported Releases 10.3.0E or later

ipv6 ospf area

Attaches an interface to an OSPF area.

Syntax `ipv6 ospf process-id area area-id`

Parameters

- *process-id*—Enter an OSPFv3 process ID for a specific OSPFv3 process, from 1 to 65535.
- `area-id`—Enter the OSPFv3 area ID in dotted decimal A.B.C.D format or enter an area ID number, from 1 to 65535.

Default Not configured

Command Mode INTERFACE

Usage Information The `no` version of this command removes an interface from an OSPFv3 area.

Example

```
OS10(config)# interface vlan 10
OS10(conf-if-vl-10)# ipv6 ospf 10 area 1
```

Supported Releases 10.3.0E or later

ipv6 ospf authentication

Configures OSPFv3 authentication on an IPv6 interface.

Syntax `ipv6 ospf authentication {null | ipsec spi number {MD5 | SHA1} key}`

Parameters

- `null` — Prevents area authentication from being inherited on the interface.
- `ipsec spi number` — Enter a unique security policy index number, from 256 to 4294967295.
- `md5` — Enable MD5 authentication.
- `sha1` — Enable SHA1 authentication.
- *key* — Enter the text string used by the authentication type.

Default IPv6 OSPF authentication is not configured on an interface.

Command Mode INTERFACE

Usage Information

- Before you enable IPsec authentication on an OSPFv3 interface, you must enable IPv6 unicast routing globally, configure an IPv6 address and enable OSPFv3 on the interface, and assign it to an area.
- The SPI value must be unique to one IPsec authentication or encryption security policy on the router. You cannot configure the same SPI value on another interface even if it uses the same authentication or encryption algorithm.
- You cannot use an IPsec MD5 or SHA1 authentication type and the `null` setting at same time on an interface. These settings are mutually exclusive.
- All neighboring OSPFv3 routers must share the key to exchange information. Only a non-encrypted key is supported. For MD5 authentication, the non-encrypted key must be 32 plain hex digits. For SHA1 authentication, the non-encrypted key must be 40 hex digits. An encrypted key is not supported.

Example

```
OS10(config)# interface ethernet 1/1/6
OS10(conf-if-eth1/1/6)# ipv6 ospf authentication ipsec spi 400 md5
12345678123456781234567812345678
```

Supported Releases 10.4.0E(R1) or later

ipv6 ospf cost

Changes the cost associated with the OSPFv3 traffic on an interface

Syntax `ipv6 ospf cost cost`

Parameters *cost* — Enter a value as the OSPFv3 cost for the interface, from 1 to 65335.

Default Based on bandwidth reference

Command Mode INTERFACE

Usage Information If not configured, the interface cost is based on the `auto-cost` command. This command configures OSPFv3 over multiple vendors to ensure that all routers use the same cost value. The `no` version of this command removes the IPv6 OSPF cost configuration.

Example

```
OS10(config)# interface vlan 10
OS10(conf-if-vl-10)# ipv6 ospf cost 10
```

Supported Releases 10.3.0E or later

ipv6 ospf dead-interval

Sets the time interval since the last hello-packet was received from a router. After the interval elapses, the neighboring routers declare the router dead.

Syntax `ipv6 ospf dead-interval seconds`

Parameters *seconds* — Enter the dead interval value in seconds, from 1 to 65535.

Default 40 seconds

Command Mode INTERFACE

Usage Information The dead interval is four times the default hello-interval by default. The `no` version of this command resets the value to the default.

Example

```
OS10(config)# interface vlan 10
OS10(conf-if-vl-10)# ipv6 ospf dead-interval 10
```

Supported Releases 10.3.0E or later

ipv6 ospf encryption

Configures OSPFv3 encryption on an IPv6 interface.

Syntax `ipv6 ospf encryption {ipsec spi number esp encryption-type key authentication-type key | null}`

Parameters

- `ipsec spi number` — Enter a unique security policy index number, from 256 to 4294967295.
- `esp encryption-type` — Enter the encryption algorithm used with ESP (3DES, DES, AES-CBC, or NULL). For AES-CBC, only the AES-128 and AES-192 ciphers are supported.
- *key* — Enter the text string used in the encryption algorithm.

- *authentication-type* — Enter the encryption MD5 or SHA1 authentication algorithm to use.
- *key* — Enter the text string used in the authentication algorithm.
- null — Enter the keyword to not use the IPsec encryption.

Default IPv6 OSPF encryption is not configured on an interface.

Command Mode INTERFACE

Usage Information

- Before you enable IPsec authentication on an OSPFv3 interface, you must enable IPv6 unicast routing globally, configure an IPv6 address and enable OSPFv3 on the interface, and assign it to an area.
- When you configure encryption on an interface, both IPsec encryption and authentication are enabled. You cannot configure encryption if you have already configured an interface for IPsec authentication using the `ipv6 ospf authentication ipsec` command. To configure encryption, you must first delete the authentication policy.
- All neighboring OSPFv3 routers must share the same encryption key to decrypt information. Only a non-encrypted key is supported. Required lengths of the non-encrypted key are: 3DES — 48 hex digits; DES — 16 hex digits; AES-CBC — 32 hex digits for AES-128 and 48 hex digits for AES-192.
- All neighboring OSPFv3 routers must share the same authentication key to exchange information. Only a non-encrypted key is supported. For MD5 authentication, the non-encrypted key must be 32 plain hex digits. For SHA1 authentication, the non-encrypted key must be 40 hex digits. An encrypted key is not supported.

Example

```
OS10(config)# interface ethernet 1/1/6
OS10(conf-if-eth1/1/6)# ipv6 ospf encryption ipsec spi 500 esp des
1234567812345678 md5 12345678123456781234567812345678
```

```
OS10(config)# interface ethernet 1/1/5
OS10(conf-if-eth1/1/5)# ipv6 ospf encryption null
```

Supported Releases 10.4.0E(R1) or later

ipv6 ospf hello-interval

Sets the time interval between hello packets sent on an interface.

Syntax `ipv6 ospf hello-interval seconds`

Parameters *seconds* — Enter the hello-interval value in seconds, from 1 to 65535.

Default 10 seconds

Command Mode INTERFACE

Usage Information All routers in a network must have the same hello time interval between the hello packets. The `no` version of the this command resets the value to the default.

Example

```
OS10(config)# interface vlan 10
OS10(conf-if-vl-10)# ipv6 ospf hello-interval 30
```

Supported Releases 10.3.0E or later

ipv6 ospf network

Sets the network type for the interface.

Syntax `ipv6 ospf network {point-to-point | broadcast}`

Parameters

- `point-to-point` — Sets the interface as part of a point-to-point network.

- `broadcast` — Sets the interface as part of a broadcast network.

Default	Broadcast
Command Mode	INTERFACE
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(config)# interface ethernet 1/1/1` `OS10(conf-if-eth1/1/1)# ipv6 ospf network broadcast`
Supported Releases	10.3.0E or later

ipv6 ospf passive

Configures an interface as a passive interface and suppresses both receiving and sending routing updates to the passive interface.

Syntax	`ipv6 ospf passive`
Parameters	None
Default	Not configured
Command Mode	INTERFACE
Usage Information	You must configure the interface before setting the interface to passive mode. The `no` version of the this command disables Passive interface configuration. **NOTE: As loopback interfaces are implicitly passive, the configuration to suppress sending and receiving of OSPF routing updates does not take effect on the loopback interfaces. However, network information corresponding to these loopback interfaces is still announced in OSPF LSAs that are sent through other interfaces configured for OSPF.**
Example	`OS10(config)# interface ethernet 1/1/6` `OS10(conf-if-eth1/1/6)# ipv6 ospf passive`
Supported Releases	10.3.0E or later

ipv6 ospf priority

Sets the priority of the interface to determine the DR for the OSPFv3 network.

Syntax	`ipv6 ospf priority` *number*
Parameters	*number* — Enter a router priority number, from 0 to 255.
Default	1
Command Mode	INTERFACE
Usage Information	When two routers attached to a network attempt to become the DR, the one with the higher router priority takes precedence. The `no` version of this command resets the value to the default.
Example	`OS10(config)# interface ethernet 1/1/6` `OS10(conf-if-eth1/1/6)# ipv6 ospf priority 4`
Supported Releases	10.3.0E or later

log-adjacency-changes

Enables logging of syslog messages about changes in the OSPFv3 adjacency state.

Syntax	`log-adjacency-changes`
Parameters	None
Default	Disabled
Command Mode	ROUTER-OSPFv3
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(config)# router ospfv3 100` `OS10(config-router-ospfv3-100)# log-adjacency-changes`
Supported Releases	10.3.0E or later

maximum-paths

Enables forwarding of packets over multiple paths.

Syntax	`maximum–paths` *`number`*
Parameters	*`number`* —Enter the number of paths for OSPFv3, from 1 to 128.
Default	Disabled
Command Mode	ROUTER-OSPFv3
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(config)# router ospfv3` `OS10(config-router-ospfv3-100)# maximum-paths 1`
Supported Releases	10.3.0E or later

redistribute

Redistributes information from another routing protocol or routing instance to the OSPFv3 process.

Syntax	`redistribute {bgp` *`as-number`* `| connected | static} [route-map` *`route-map name`*`]`
Parameters	• *`as-number`* — Enter an autonomous number to redistribute BGP routing information throughout the OSPFv3 instance, from 1 to 4294967295. • *`route-map name`* — Enter the name of a configured route-map. • `connected` — Enter the information from the connected active routes on interfaces to redistribute. • `static` — Enter the information from static routes on interfaces redistribute.
Defaults	Not configured
Command Mode	ROUTER-OSPFv3
Usage Information	When an OSPFv3 redistributes, the process is not completely removed from the BGP configuration. The `no` version of this command disables the redistribute configuration.
Example	`OS10(config)# router ospfv3 100` `OS10(config-router-ospfv3-100)# redistribute bgp 4 route-map dell1`

Example (Connected)	`OS10((config-router-ospfv3-100)# redistribute connected route-map dell2`
Supported Releases	10.3.0E or later

router-id

Configures a fixed router ID for the OSPFv3 process.

Syntax	`router-id` *ip-address*
Parameters	*ip-address* — Enter the IP address of the router as the router ID.
Default	Not configured
Command Mode	ROUTER-OSPFv3
Usage Information	Configure an arbitrary value in the IP address format for each router. Each router ID must be unique. Use the fixed router ID for the active OSPFv3 router process. Changing the router ID brings down the existing OSPFv3 adjacency. The new router ID becomes effective immediatley. The `no` version of this command disables the router ID configuration.
Example	`OS10(config)# router ospfv3 10` `OS10(config-router-ospfv3-100)# router-id 10.10.1.5`
Supported Releases	10.3.0E or later

router ospfv3

Enters Router OSPFv3 mode and configures an OSPFv3 instance.

Syntax	`router ospfv3` *instance-number* `[vrf` *vrf-name*`]`
Parameters	• *instance-number*—Enter a router OSPFv3 instance number, from 1 to 65535. • `vrf` *vrf-name* — Enter the keyword `vrf` followed by the name of the VRF to configure an OSPFv3 instance in that VRF.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command deletes an OSPFv3 instance.
Example	`OS10(config)# router ospfv3 10 vrf vrf-test`
Supported Releases	10.3.0E or later

show ipv6 ospf

Displays OSPFv3 instance configuration information.

Syntax	`show ipv6 ospf [`*instance-number*`]`
Parameters	*instance-number* — (Optional) View OSPFv3 information for a specified instance number, from 1 to 65535.
Default	None
Command Mode	EXEC
Usage Information	None

Example

```
OS10# show ipv6 ospf
Routing Process ospfv3 200 with ID 1.1.1.1
It is an Area Border Router
Min LSA origination 5000 msec, Min LSA arrival 1000 msec
Min LSA hold time 0 msec, Max LSA wait time 0 msec
Number of area in this router is 2, normal 2 stub 0 nssa
   Area  (0.0.0.0)
       Number of interface in this area is 1
       SPF algorithm executed 42 times
   Area  (0.0.0.1)
       Number of interface in this area is 1
       SPF algorithm executed 42 times
OS10# show ipv6 ospf 200
Routing Process ospfv3 200 with ID 10.0.0.2
Min LSA origination 5000 msec, Min LSA arrival 1000 msec
Min LSA hold time 0 msec, Max LSA wait time 0 msec
Number of area in this router is 1, normal 1 stub 0 nssa
   Area  (0.0.0.0)
       Number of interface in this area is 1
       SPF algorithm executed 3 times
```

Supported Releases 10.3.0E or later

show ipv6 ospf database

Displays all LSA information. You must enable OSPFv3 to generate output.

Syntax `show ipv6 ospf process-id [vrf vrf-name] database`

Parameters

- *process-id* — Enter the OSPFv3 process ID to view a specific process. If you do not enter a process ID, the command applies to all the configured OSPFv3 processes.
- `vrf` *vrf-name* — Enter the keyword `vrf` followed by the name of the VRF to display LSA information for that VRF.

Default Not configured

Command Mode EXEC

Usage Information

- `Link ID`—Identifies the router ID.
- `ADV Router`—Identifies the advertising router's ID.
- `Age`—Displays the LS age.
- `Seq#`—Identifies the LS sequence number. This identifies old or duplicate LSAs.
- `Checksum`—Displays the Fletcher checksum of an LSA's complete contents.
- `Link count`—Displays the number of interfaces for that router.
- `Rtr Count`—Displays the router count.
- `Dest RtrID`—Displays the destination router ID.
- `Interface`—Displays the interface type.
- `Prefix`—Displays the prefix details.

Example

```
OS10# show ipv6 ospf database
       OSPF Router with ID (10.0.0.2) (Process ID 200)
Router Link States (Area 0.0.0.0)
ADV Router      Age       Seq#        Fragment ID Link count Bits
-----------------------------------------------------------------
1.1.1.1         1610      0x80000144  0           1          B
2.2.2.2         1040      0x8000013A  0           1
10.0.0.2        1039      0x80000002  0           1
Net Link States (Area 0.0.0.0)
```

```
ADV Router      Age       Seq#        Link ID   Rtr count
-------------------------------------------------------
2.2.2.2         1045      0x80000001  5         2
Inter Area Router States (Area 0.0.0.0)
ADV Router      Age       Seq#        Link ID   Dest RtrID
-------------------------------------------------------------
1.1.1.1         1605      0x80000027  1         3.3.3.3
Link (Type-8) Link States (Area 0.0.0.0)
ADV Router      Age       Seq#        Link ID   Interface
-------------------------------------------------------------------
1.1.1.1         1615      0x80000125  5         ethernet1/1/1
2.2.2.2         1369      0x8000011B  5         ethernet1/1/1
10.0.0.2        1044      0x80000001  5         ethernet1/1/1
Type-5 AS External Link States
ADV Router      Age       Seq#        Prefix
-----------------------------------------------------------------------
3.3.3.3         3116      0x80000126  400::/64
3.3.3.3         3116      0x80000124  34::/64
```

Supported Releases 10.3.0E or later

show ipv6 ospf interface

Displays the configured OSPFv3 interfaces. You must enable OSPFv3 to display the output.

Syntax `show ipv6 ospf interface` *`interface`* `[vrf` *`vrf-name`*`]`

Parameters

- *`interface`* — (Optional) Enter the interface information:
  - `ethernet` — Physical interface, from 1 to 48.
  - `port-channel` — Port-channel interface, from 1 to 128.
  - `vlan` — VLAN interface, from 1 to 4093.
- `vrf` *`vrf-name`* — (Optional) Enter the keyword `vrf` followed by the name of the VRF to display the configured OSPFv3 enabled interfaces in that VRF.

Default Not configured

Command Mode EXEC

Example

```
OS10# show ipv6 ospf interface
ethernet1/1/1 is up, line protocol is up
  Link Local Address fe80::20c:29ff:fe0a:d59/64, Interface ID 5
  Area 0.0.0.0, Process ID 200, Instance ID 0, Router ID 10.0.0.2
  Network Type broadcast, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1
  BFD enabled(Interface level) Interval 300 Min_rx 300 Multiplier 3 Role Active
  Designated Router on this network is 2.2.2.2
  Backup Designated router on this network is 10.0.0.2 (local)
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 2.2.2.2(Designated Router)
```

Supported Releases 10.3.0E or later

show ipv6 ospf neighbor

Displays a list of OSPFv3 neighbors connected to the local router.

Syntax `show ipv6 ospf [vrf` *`vrf-name`*`] neighbor`

Parameters

- `vrf vrf-name` — Enter the keyword `vrf` followed by the name of the VRF to display a list of OSPFv3 neighbors in that VRF.

Default Not configured

Command Mode EXEC

Usage Information

- `Neighbor ID`—Displays the neighbor router ID.
- `Pri`—Displays the priority assigned neighbor.
- `State`—Displays the OSPF state of the neighbor.
- `Dead Time`—Displays the expected time until the system declares the neighbor dead.
- `Interface ID`—Displays the neighbor interface ID
- `Interface`—Displays the interface type, node/slot/port or number information.

Example

```
OS10(conf-if-eth1/1/1)# show ipv6 ospf neighbor
Neighbor ID    Pri     State     Dead Time    Interface ID  Interface
-------------------------------------------------------------------
2.2.2.2        1       Full/DR   00:00:30     5             ethernet1/1/1
```

Supported Releases 10.3.0E or later

show ipv6 ospf statistics

Displays OSPFv3 traffic statistics.

Syntax `show ipv6 ospf [instance-number] statistics [interface interface]`

Parameters

- *instance-number* — (Optional) Enter an OSPFv3 instance number, from 1 to 65535.
- `interface interface` — (Optional) Enter the interface information:
  - `ethernet node/slot/port[:subport]` — Enter an Ethernet port interface.
  - `port-channel number` — Enter the port-channel interface number, from 1 to 128.
  - `vlan vlan-id` — Enter the VLAN ID number, from 1 to 4093.

Default Not configured

Command Mode EXEC

Usage Information This command displays OSPFv3 traffic statistics for a specified instance or interface, or for all OSPFv3 instances and interfaces.

Example

```
OS10# show ipv6 ospf interface ethernet 1/1/1

Interface ethernet1/1/1
  Receive Statistics
    rx-invalid 0 rx-invalid-bytes 0
    rx-hello 0 rx-hello-bytes 0
    rx-db-des 0 rx-db-des-bytes 0
    rx-ls-req 0 rx-ls-req-bytes 0
    rx-ls-upd 0 rx-ls-upd-bytes 0
    rx-ls-ack 0 rx-ls-ack-bytes 0
  Transmit Statistics
    tx-hello 1054 tx-hello-bytes 37944
    tx-db-des 0 tx-db-des-bytes 0
    tx-ls-req 0 tx-ls-req-bytes 0
    tx-ls-upd 0 tx-ls-upd-bytes 0
    tx-ls-ack 0 tx-ls-ack-bytes 0
  Error packets (Receive statistics)
```

```
    bad-src 0 dupe-id 0 hello-err 0
    mtu-mismatch 0 nbr-ignored 0
    resource-err 0 bad-lsa-len 0 lsa-bad-type 0
    lsa-bad-len 0 lsa-bad-cksum 0
    hello-tmr-mismatch 0 dead-ivl-mismatch 0
    options-mismatch 0 nbr-admin-down 0 own-hello-drop 0
    self-orig 0 wrong-length 0
    version-mismatch 0 area-mismatch 0
```

Supported Releases 10.4.0E(R1) or later

timers spf (OSPFv3)

Enables shortest path first (SPF) throttling to delay an SPF calculation when a topology change occurs.

Syntax `timers spf [start-time [hold-time [max-wait]]]`

Parameters

- *start-time* — Sets the initial SPF delay in milliseconds, from 1 to 600000; default 1000.
- *hold-time* — Sets the additional hold time between two SPF calculations in milliseconds, from 1 to 600000; default 10000.
- *max-wait* — Sets the maximum wait time between two SPF calculations in milliseconds, from 1 to 600000; default 10000.

Default

- *start-time* — 1000 milliseconds
- *hold-time* — 10000 milliseconds
- *max-wait* — 10000 milliseconds

Command Mode ROUTER-OSPFv3

Usage Information OSPFv2 and OSPFv3 support SPF throttling. By default, SPF timers are disabled in an OSPF instance.

Use SPF throttling to delay SPF calculations during periods of network instability. In an OSPF network, a topology change event triggers an SPF calculation after a specified start time. When the start timer finishes, a hold time may delay the next SPF calculation for an additional time. When the hold timer is running:

- Each time a topology change occurs, the SPF calculation delays for double the configured hold time up to maximum wait time.
- If no topology change occur, an SPF calculation performs and the hold timer resets to its configured value.

If you do not specify a start-time, hold-time, or max-wait value, the default values are used. The `no` version of this command removes the configured SPF timers and disables SPF throttling in an OSPF instance.

Example

```
OS10(config)# router ospfv3 100
OS10(config-router-ospfv3-100)# timers spf 1345 2324 9234

OS10(config-router-ospfv3-100)# do show ipv6 ospf
Routing Process ospfv3 100 with ID 129.240.244.107
SPF schedule delay 1345 msecs, Hold time between two SPFs 2324 msecs
Min LSA origination 5000 msec, Min LSA arrival 1000 msec
Min LSA hold time 0 msec, Max LSA wait time 0 msec
Number of area in this router is 1, normal 1 stub 0 nssa
Area (0.0.0.1)
Number of interface in this area is 1
SPF algorithm executed 2 times
```

Supported Releases 10.4.0E(R1) or later

Object tracking manager

OTM allows you to track the link status of Layer 2 (L2) interfaces, and the reachability of IPv4 and IPv6 hosts. You can increase the availability of the network and shorten recovery time if an object state goes Down.

Object tracking monitors the status of tracked objects and communicates any changes made to interested client applications. OTM client applications are virtual router redundancy protocol (VRRP) and policy-based routing (PBR). Each tracked object has a unique identifying number that clients use to configure the action to take when a tracked object changes state. You can also optionally specify a time delay before changes in a tracked object's state report to a client application.

VRRP subscribes to a track object which tracks the interface line protocol state. It uses the tracked object status to determine the priority of the VRRP router in a VRRP group. If a tracked state or interface goes down, VRRP updates the priority based on how you configure the new priority for the tracked state. When the tracked state comes up, VRRP restores the original priority for the virtual router group.

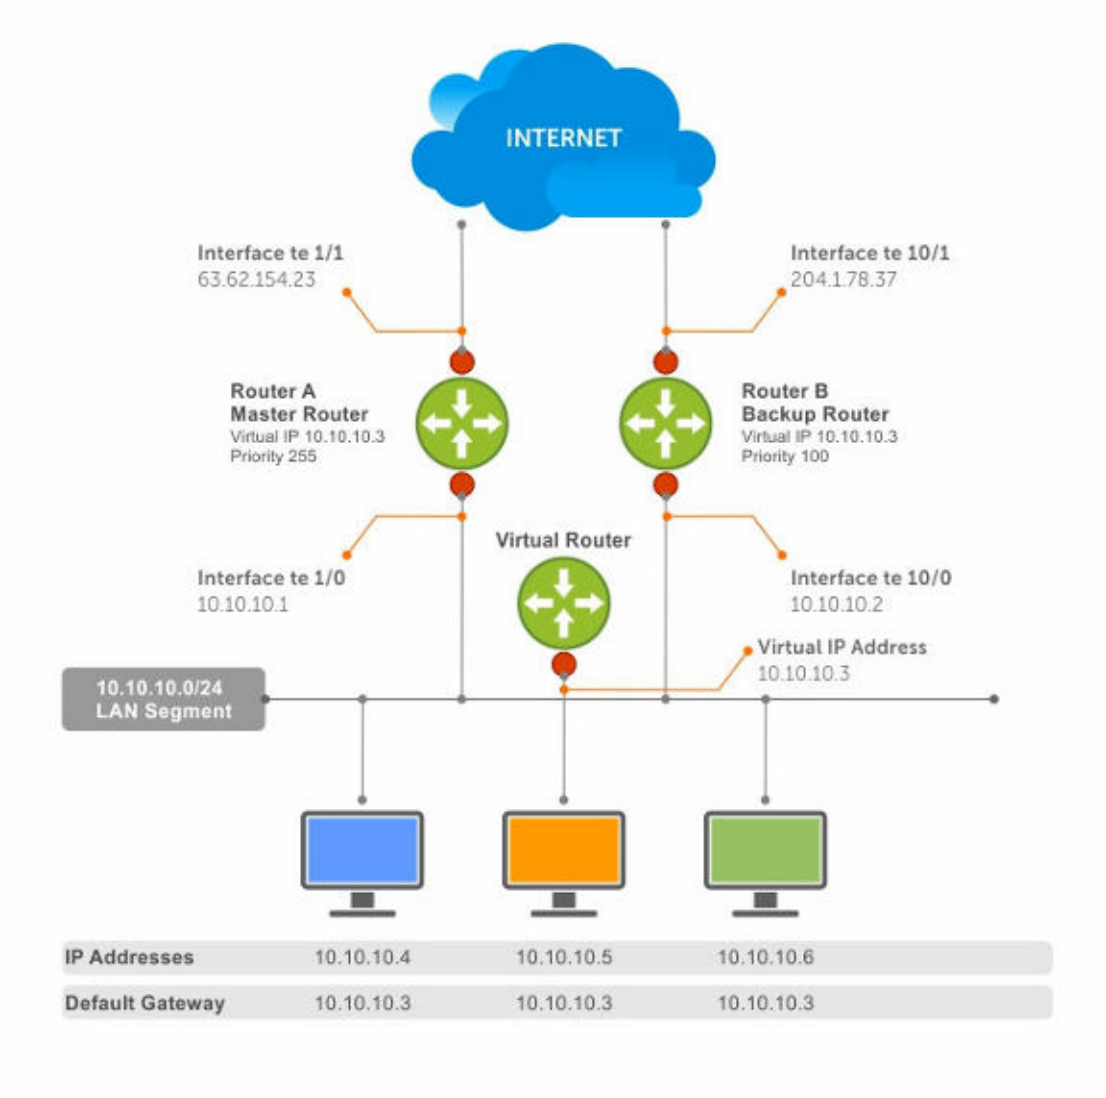


Figure 4. Object tracking

Interface tracking

You can create an object that tracks the line-protocol state of an L2 interface, and monitors its operational up or down status. You can configure up to 500 objects. Each object is assigned a unique ID.

When the link-level status goes down, the tracked resource status is also considered Down. If the link-level status goes up, the tracked resource status is also considered Up. For logical interfaces such as port-channels or VLANs, the link-protocol status is considered Up if any physical interface under the logical interface is Up.

The list of available interfaces include:

- `ethernet` — Physical interface
- `port-channel` — Port-channel identifier
- `VLAN` — Virtual local area network (VLAN) identifer
- `Loopback` — Loopback interface identifier
- `mgmt` — Management interface

1 Configure object tracking in CONFIGURATION mode, from 1 to 500.
```
track object-id
```
2 (Optional) Enter interface object tracking on the line-protocol state of an L2 interface in OBJECT TRACKING mode.
```
interface interface line-protocol
```
3 (Optional) Configure the time delay used before communicating a change to the status of a tracked interface in OBJECT TRACKING mode, from 0 to 80 seconds; default 0.
```
delay [up seconds] [down seconds]
```
4 (Optional) View the tracked object information in EXEC mode.
```
show track object-id
```
5 (Optional) View all interface object information in EXEC mode.
```
show track interface
```
6 (Optional) View all IPv4 or IPv6 next-hop object information in EXEC mode.
```
show track [ip | ipv6]
```
7 (Optional) View brief status of object information in EXEC mode.
```
show track brief
```

Configure object tracking

```
OS10(config)# track 1
OS10(conf-track-1)# interface ethernet 1/1/1 line-protocol
OS10(conf-track-1)# delay up 20
OS10(conf-track-1)# delay down 10
OS10(conf-track-1)# do show track 1
Interface ethernet1/1/1 line-protocol
Line protocol is UP
1 changes, Last change 2017-04-26T06:41:36Z
```

Host tracking

If you configure an IP host as a tracked object, the entry or next-hop address in the ARP cache determines the Up or Down state of the route.

A tracked host is reachable if there is an ARP cache entry for the router's next-hop address. An attempt to regenerate the ARP cache entry occurs if the next-hop address appears before considering the route Down.

1 Configure object tracking in CONFIGURATION mode.

```
track object-id
```

2 Enter the host IP address for reachability of an IPv4 or IPv6 route in OBJECT TRACKING mode.

```
[ip | ipv6] host-ip-address reachability
```

3 Configure the time delay used before communicating a change in the status of a tracked route in OBJECT TRACKING mode.

```
delay [up seconds] [down seconds]
```

4 Track the host by checking the reachability periodically in OBJECT TRACKING mode.

```
reachability-refresh interval
```

5 View the tracking configuration and the tracked object status in EXEC mode.

```
show track object-id
```

Configure IPv4 host tracking

```
OS10 (conf-track-1)# track 2
OS10 (conf-track-2)# ip 1.1.1.1 reachability
OS10 (conf-track-2)# do show track 2
IP Host 1.1.1.1 reachability
Reachability is DOWN
1 changes, Last change 2017-04-26T06:45:31Z
OS10 (conf-track-2)#
```

Configure IPv6 host tracking

```
OS10 (conf-track-2)# track 3
OS10 (conf-track-3)# ipv6 20::20 reachability
OS10 (conf-track-3)# delay up 20
OS10 (conf-track-3)# do show track 3
IP Host 20::20 reachability
Reachability is DOWN
1 changes, Last change 2017-04-26T06:47:04Z
OS10 (conf-track-3)#
```

Set tracking delays

You can configure an optional Up or Down timer for each tracked object. The timer allows you to set the time delay before a change in the state of a tracked object communicates to the clients. The time delay starts when the state changes from Up to Down or from Down to Up.

If the state of an object changes back to its former Up or Down state before the timer expires, the timer is canceled without notifying the client. If the timer expires and an object's state has changed, a notification is sent to the client. For example, if the Down timer is running and an interface goes down then comes back up, the Down timer is canceled. The client is not notified of the event.

If you do not configure a delay, a notification is sent when a change in the state of a tracked object is detected. The time delay in communicating a state change is specified in seconds.

Object tracking

As a client, VRRP can track up to 20 interface objects plus 12 tracked interfaces supported for each VRRP group. You can assign a unique priority-cost value, from 1 to 254, to each tracked VRRP object or group interface.

If a tracked VRRP object is in a Down state, the priority cost is subtracted from the VRRP group priority. If a VRRP group router acts as owner-master, the run-time VRRP group priority remains fixed at 255. Changes in the state of a tracked object have no effect.

In VRRP object tracking, the sum of the priority costs for all tracked objects and interfaces cannot equal or exceed the priority of the VRRP group.

View tracked objects

You can view the status of currently tracked L2 or L3 interfaces, or the IPv4 or IPv6 hosts.

View brief object tracking information

```
OS10# show track brief
TrackID   Resource               Parameter           Status        LastChange
-------------------------------------------------------------------------------
1         line-protocol          ethernet1/1/1       DOWN           2017-02-03T08:41:25Z1
2         ipv4-reachablity       1.1.1.1             DOWN           2017-02-03T08:41:43Z1
3         ipv6-reachablity       10::10              DOWN           2017-02-03T08:41:55Z1
```

View all object tracking information

```
OS10# show track
```

View interface object tracking information

```
OS10# show track interface
TrackID   Resource               Parameter           Status        LastChange
-------------------------------------------------------------------------------
1         line-protocol          ethernet1/1/1       DOWN           2017-02-03T08:41:25Z1
OS10# show track ip
TrackID   Resource               Parameter           Status        LastChange
-------------------------------------------------------------------------------
2         ipv4-reachablity       1.1.1.1             DOWN           2017-02-03T08:41:43Z1
OS10# show track ipv6
TrackID   Resource               Parameter           Status        LastChange
-------------------------------------------------------------------------------
3         ipv6-reachablity       10::10              DOWN           2017-02-03T08:41:55Z1
```

View IPv4 next-hop object tracking

```
OS10# show track ip
```

View IPv6 next-hop object tracking

```
OS10# show track ipv6
```

View running configuration

```
OS10# show running-configuration
```

OTM commands

delay

Configures the delay timers.

Syntax	`delay {up | down} seconds`
Parameters	*seconds* — Enter the delay time in seconds. A maximum of 180 characters.
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	None

Example	`OS10(conf-track-100)# delay up 200 down 100`
Supported Releases	10.3.0E or later

interface line-protocol

Configures an object to track a specific interface's line-protocol status.

Syntax	`interface` *`interface`* `line-protocol`
Parameters	*`interface`* — Enter the interface information: • `ethernet` — Physical interface. • `port-channel` — Enter the port-channel identifier. • `vlan` — Enter the VLAN identifier. • `loopback` — Enter the Loopback interface identifier. • `mgmt` — Enter the Management interface.
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	None
Example	`OS10(conf-track-100)# interface ethernet line-protocol`
Supported Releases	10.3.0E or later

ip reachability

Configures an object to track a specific next-hop host's reachability.

Syntax	`ip` *`host-ip-address`* `reachability`
Parameters	*`host-ip-address`* — Enter the IPv4 host address.
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	None
Example	`OS10(config)# track 100` `OS10(conf-track-100)# ip 10.10.10.1 reachability`
Supported Releases	10.3.0E or later

ipv6 reachability

Configures an object to track a specific next-hop host's reachability.

Syntax	`ipv6` *`host-ip-address`* `reachability`
Parameters	*`host-ip-address`* — Enter the IPv6 host address.
Defaults	Not configured

Command Mode	CONFIGURATION
Usage Information	None
Example	`OS10(config)# track 200` `OS10(conf-track-200)# ipv6 10::1 reachability`
Supported Releases	10.3.0E or later

reachability-refresh

Configures a polling interval for reachability tracking.

Syntax	`reachability-refresh` *`interval`*
Parameters	*`interval`* — Enter the polling interval value. A maximum of 3600 seconds.
Defaults	0 seconds
Command Mode	CONFIGURATION
Usage Information	Set the interval to 0 to disable the refresh.
Example	`OS10(conf-track-100)# reachability-refresh 600`
Supported Releases	10.3.0E or later

show track

Displays tracked object information.

Syntax	`show track [`*`brief`*`] [`*`object-id`*`] [interface] [ip | ipv6]`
Parameters	• *`brief`* — (Optional) Displays brief tracked object information. • *`object-id`* — (Optional) Displays tracked object information for a specific object ID. • `interface` — (Optional) Displays all interface object information. • `ip` — (Optional) Displays all IPv4 next-hop object information. • `ipv6` — (Optional) Displays all IPv6 next-hop object information.
Defaults	None
Command Mode	CONFIGURATION
Usage Information	None

Example (Brief)	`OS10# show track brief`
	`TrackID  Resource            Parameter       Status  LastChange`
	`----------------------------------------------------------------------`
	`1        line-protocol       ethernet1/1/1   DOWN    2017-02-03T08:41:25Z1`
	`2        ipv4-reachablity    1.1.1.1         DOWN    2017-02-03T08:41:43Z1`
	`3        ipv6-reachablity    10::10          DOWN    2017-02-03T08:41:55Z1`

Supported Releases	10.3.0E or later

track

Configures and manages tracked objects.

Syntax	`track object-id`
Parameters	*object-id* — Enter the object ID to track. A maximum of 500.
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command deletes the tracked object from an interface.
Example	`OS10# track 100`
Supported Releases	10.3.0E or later

Policy-based routing

PBR provides a mechanism to redirect IPv4 and IPv6 data packets based on the policies defined to override the switch's forwarding decisions based on the routing table.

Policy-based route-maps

A route-map is an ordered set of rules that controls the redistribution of IP routes into a protocol domain. When you enable PBR on an interface, all IPv4 or IPv6 data packets process based on the policies that you define in the route-maps. The rules defined in route-maps are based on access control lists (ACLs) and next-hop addresses, and only apply to ACLs used in policy-based routing.

You can create a route-map that specifies the match criteria and resulting action if all the match clauses are met. After you create the route-map, you can enable PBR for that route-map on a specific interface. Route-maps contain `match` and `set` statements that you can mark as *permit*.

Access-list to match route-map

You can assign an IPv4 or IPv6 access-list to match a route-map. The IP access list contains the criteria to match the traffic content based on the header field, such as the destination IP or source IP.

When `permit` or `deny` is present in the `access-list`, it is omitted and the action present in the `route-map` command is used for policy-based routing. The `permit` keyword in the route-map statement indicates policy-based routing. The `deny` keyword in the route-map statement indicates a switch-based forwarding decision, a PBR exception. Only use access list for the packet match criteria in policy-based routing.

1. Assign an access-list to match the route-map in CONFIGURATION mode.
   `ip access-list access-list-name`
2. Set the IP address to match the access-list in IP-ACL mode.
   `permit ip ip-address`

Configure IPv4 access-list to match route-map

```
OS10(config)# ip access-list acl5
OS10(conf-ipv4-acl)# permit ip 10.10.10.0/24 any
```

Configure IPv6 access-list to match route-map

```
OS10(config)# ipv6 access-list acl8
OS10(conf-ipv6-acl)# permit ipv6 10::10 any
```

Set address to match route-map

You can set an IPv4 or IPv6 address to match a route-map.

1 Enter the IPv4 or IPv6 address to match and specify the access-list name in Route-Map mode.

```
match {ip | ipv6} address access-list-name
```

2 Set the next-hop IP address in Route-Map mode.

```
set {ip | ipv6} next-hop ip-address
```

Apply match and set parameters to IPv4 route-map

```
OS10(conf-route-map)# route-map map1
OS10(conf-route-map)# match ip address acl5
OS10(conf-route-map)# set ip next-hop 10.10.10.10
```

Apply match and set parameters to IPv6 route-map

```
OS10(conf-route-map)# route-map map1
OS10(conf-route-map)# match ipv6 address acl8
OS10(conf-route-map)# set ipv6 next-hop 20::20
```

Assign route-map to interface

You can assign a route-map to an interface for IPv4 or IPv6 policy-based routing to an interface.

- Assign the IPv4 or IPv6 policy-based route-map to an interface in INTERFACE mode.

```
{ip | ipv6} policy route-map map-name
```

Assign route-map to an IPv4 interface

```
OS10(config)# interface ethernet 1/1/5
OS10(conf-if-eth1/1/5)# ip policy route-map map1
```

Assign route-map to an IPv6 interface

```
OS10(conf-if-eth1/1/5)# ipv6 policy route-map map2
```

View PBR information

Display PBR information to verify IPv4 or IPv6 configuration and view statistics.

1 View IPv4 or IPv6 PBR policy information in EXEC mode.

```
show {ip | ipv6} policy name
```

2 View current PBR statistics in EXEC mode.

```
show route-map map-name pbr-statistics
```

3 Clear all policy statistics information in EXEC mode.

```
clear route-map map-name pbr-statistics
```

Verify IPv4 PBR configuration

```
OS10# show ip policy abc
Interface      Route-map
-----------------------
```

```
ethernet1/1/1   abc
ethernet1/1/3   abc
vlan100         abc
```

Verify IPv6 PBR configuration

```
OS10# show ipv6 policy abc
Interface       Route-map
------------------------
ethernet1/1/1   abc
ethernet1/1/3   abc
vlan100         abc
```

View current PBR statistics

```
show route-map pbr-sample pbr-statistics
route-map pbr-sample, permit, sequence 10

Policy routing matches: 84 packets
```

Policy-based routing per VRF

Configure PBR per VRF instance for both IPv4 and IPv6 traffic flows.

Policy-based routing (PBR) enables packets with certain match criteria, such as packets from specific source and destination addresses, to be re-directed to a different next-hop.

You can also use PBR to re-direct packets arriving on a VRF instance to a next-hop that is reachable through a different VRF instance. You can re-direct packets arriving on any VRF instance to the default VRF instance or any other non-default VRF instance.

Configuring PBR per VRF

For traffic arriving on a VRF instance, you can re-direct this traffic to a next-hop on another VRF instance using route-maps. In the route-map, set the next-hop IP address that is reachable through a different VRF instance. When traffic that matches certain criteria arrives on a VRF instance, the route-map configuration enables packets to be re-directed to a next-hop that is reachable over another VRF instance. To configure PBR per VRF:

1 Create the match ACL rule for IPv4 or IPv6 traffic.

   `{ip | ipv6} access-list` *`access-list-name`*

2 Permit or deny IPv4 or IPv6 traffic from any source with a specific destination.

   `permit {ip | ipv6} any` *`ip-address`*

   or

   `deny {ip | ipv6} any` *`ip-address`*

3 Configure a route-map to re-direct traffic arriving on a specific VRF instance.

   `route-map route-map-name`

4 Enter the IPv4 or IPv6 address to match and specify the access-list name.

   `match {ip | ipv6} address access-list-name`

5 In the route-map, set the IPv4 or IPv6 next-hop to be reached through a different VRF instance.

   `set {ip | ipv6} vrf` *`vrf-name`* `next-hop` *`next-hop-ipv4address`*

   This next-hop-address is reachable through a different VRF instance.

   NOTE: If the next-hop is reachable on the specified VRF instance, the packet is redirected; otherwise, the packet follows the regular routing flow.

6 Apply the route-map to the interface.

```
interface interface-type

{ip | ipv6} policy route-map route-map-name
```

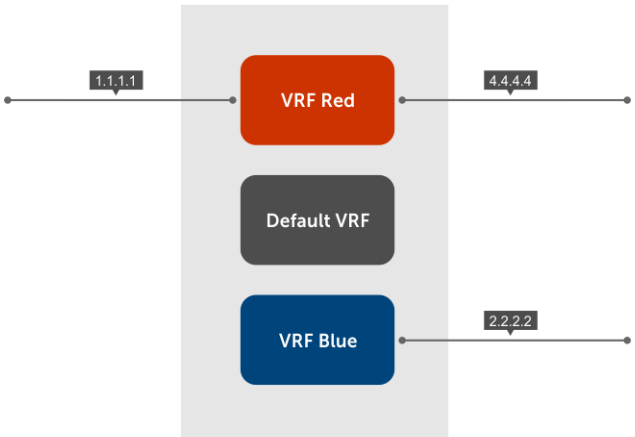
7 View the route-map information.

```
show route-map
```

```
OS10(conf-if-vl-40)# do show route-map
route-map test, permit, sequence 10
Match clauses:
ip address (access-lists): acl1
Set clauses:
ip vrf red next-hop 1.1.1.1 track-id 200
```

Sample configuration

Consider a scenario where traffic from source IP address 1.1.1.1 ingresses through VLAN40 that is part of VRF RED. The egress interface for this traffic is also on the same VRF RED with IP address 4.4.4.4, as shown.



Using the following PBR configuration, you can re-direct traffic ingresssing to VRF RED to a destination that is reachable through the next-hop IP address `2.2.2.2` in VRF BLUE:

1 Create a route-map.

```
OS10(config)# route-map test
```

2 Enter the IP address to match the specified access list.

```
OS10(config-route-map)# match ip 4.4.4.4 acl1
```

3 Set the next-hop address to `2.2.2.2`, which is reachable through VRF BLUE.

```
OS10(config-route-map)#
  OS10(config-route-map)# set ip vrf BLUE next-hop 2.2.2.2
OS10(config-route-map)# exit
```

4 Apply this rule to the interface where the traffic ingresses, in this case VLAN40.

```
OS10(config)# interface vlan 40
OS10(conf-if-vl-40)#
OS10(conf-if-vl-40)# ip policy route-map test
```

5 (Optional) View the PBR configuration on the interface.

```
OS10(conf-if-vl-40)# show configuration
 !
  interface vlan40
```

```
    no shutdown
    ip policy route-map test
   !
```

Track route reachability

Track IPv4 or IPv6 reachablility using object tracking. To configure tracking over the routes that are reachable through a VRF instance:

1 Configure object tracking.

```
track track-id
```

```
OS10(config)# track 200
```

2 Configure reachability of the next-hop address through the VRF instance.

```
ip ip-address reachablility vrf vrf-name
```

```
OS10(conf-track-200)#
OS10(conf-track-200)# ip 1.1.1.1 reachability vrf red
OS10(conf-track-200)#exit
```

3 Configure the route-map.

```
route-map route-map-name
```

```
OS10(config-route-map)#
OS10(config-route-map)# match ip address acl1
```

4 Set the track ID configured in step 1 to the route-map.

```
set ip vrf vrf-name nexy-hop next-hop-address track-id track-id-number
```

```
OS10(config-route-map)# set ip vrf red next-hop 1.1.1.1 track-id 200
```

5 Apply the route-map to the interface where traffic is ingressing on the VRF instance.

```
interface interface-type

ip policy route-map route-map-name
```

```
OS10(config)# interface vlan 40
OS10(conf-if-vl-40)#
OS10(conf-if-vl-40)# ip policy route-map test
OS10(conf-if-vl-40)# show configuration
!
```

NOTE: Ensure you configure track and PBR next-hop with the same VRF instance. For next-hop reachability in the same VRF instance, you must configure both PBR per VRF and object tracking. Missing either the track or PBR next-hop configuration in a VRF instance results in an erroneous configuration. However, the system does not display an error message indicating problems in the configuration.

View PBR configuration

Use the `show configuration` command to view the configuration of the interface.

```
OS10(conf-if-vl-40)# show configuration
!
interface vlan40
no shutdown
ip policy route-map test
```

Use the `show route-map` command to view the route-map configuration.

```
OS10(config)# do show route-map
 route-map map1, permit, sequence 10
 Match clauses:
    ipv6 address (access-lists): acl1
 Set clauses:
    ipv6 vrf {vrf-name} next-hop 5555::5556

OS10(conf-if-vl-40)# do show route-map
route-map test, permit, sequence 10
Match clauses:
ip address (access-lists): acl1
Set clauses:
ip next-hop 1.1.1.1 track-id 200

OS10(conf-if-vl-40)# do show route-map test
route-map test, permit, sequence 10
Match clauses:
ip address (access-lists): acl1
Set clauses:
ip vrf red next-hop 1.1.1.1 track-id 200
!
```

PBR commands

clear route-map pbr-statistics

Clears all PBR counters.

Syntax	`clear route-map [map-name] pbr-statistics`
Parameters	*map-name*—Enter the name of a configured route-map. A maximum of 140 characters.
Defaults	None
Command Mode	EXEC
Usage Information	None
Example	`OS10# clear route-map map1 pbr-statistics`
Supported Releases	10.3.0E or later

match address

Matches the access-list to the route-map.

Syntax	`match {ip | ipv6} address [name]`
Parameters	*name*—Enter the name of an access-list. A maximum of 140 characters.
Defaults	Not configured
Command Mode	ROUTE-MAP
Usage Information	None
Example	`OS10(conf-route-map)# match ip address acl1`

Supported Releases	10.3.0E or later

policy route-map

Assigns a route-map for IPv4 or IPV6 policy-based routing to the interface.

Syntax	`{ip | ipv6} policy route-map [map-name]`
Parameters	*map-name*—Enter the name of a configured route-map. A maximum of 140 characters.
Defaults	Not configured
Command Mode	INTERFACE
Usage Information	None
Example	`OS10(config)# interface ethernet 1/1/1` `OS10(conf-if-eth1/1/1)# ip policy route-map map1`
Supported Releases	10.3.0E or later

route-map pbr-statistics

Enables counters for PBR statistics.

Syntax	`route-map [map-name] pbr-statistics`
Parameters	*map-name*—Enter the name of a configured route-map. A maximum of 140 characters.
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	None
Example	`OS10(config)# route-map map1 pbr-statistics`
Supported Releases	10.3.0E or later

set next-hop

Sets an IPv4 or IPv6 next-hop address for policy-based routing.

Syntax	`set {ip | ipv6} vrf [vrf-name] next-hop address`
Parameters	• `vrf vrf-name` — Enter the keyword then the name of the VRF to make the next-hop reachable over that VRF. • *address* — Enter the next-hop IPv4 or IPv6 address.
Defaults	Not configured
Command Mode	ROUTE-MAP
Usage Information	You must configure track and PBR next-hop with the same VRF instance. For next-hop reachability in the same VRF instance, you must configure both PBR per VRF and object tracking. Missing either the track or PBR next-

hop configuration in a VRF instance results in an erroneous configuration. However, the system does not display an error message indicating problems in the configuration.

Example

```
OS10(conf-route-map)# set ip next-hop 10.10.10.10
*Sets the next-hop IP address.
OS10(conf-route-map)#set ip vrf red next-hop 2.2.2.2
*The next-hop 2.2.2.2 should be reachable via interface over VRF "red".
```

Supported Releases 10.3.0E or later

set next-hop track

Sets the next-hop IPv4 or IPv6 address to track the PBR object.

Syntax `set {ip | ipv6} vrf [vrf-name] next-hop address track track-id`

Parameters

- *address*—Enter an IPv4 or IPv6 address.
- *vrf vrf-name* — Enter the keyword then the name of the VRF to track the next-hop reachable through that VRF.
- *track-id*—(Optional) Enter the track ID of the PBR object.

Defaults Not configured

Command Mode ROUTE-MAP

Usage Information You must configure track and PBR next-hop with the same VRF instance. For next-hop reachability in the same VRF instance, you must configure both PBR per VRF and object tracking. Missing either the track or PBR next-hop configuration in a VRF instance results in an erroneous configuration. However, the system does not display an error message indicating problems in the configuration.

Example

```
OS10(conf-route-map)# set ip next-hop 10.10.10.10 track-id 12
*Set the track ID configured to the route-map.
OS10(conf-route-map)# set ip vrf red next-hop 1.1.1.1 track-id 200
*Sets the track ID configured to track the next-hop reachable through the VRF
specified.
```

Supported Releases 10.3.0E or later

show policy

Displays policy information.

Syntax `show {ip | ipv6} policy [map-name]`

Parameters *map-name* — (Optional) Enter the name of a configured route map. A maximum of 140 characters.

Defaults None

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip policy map-name
```

Supported Releases 10.3.0E or later

show route-map pbr-statistics

Displays the current PBR statistics.

Syntax	`show route-map [map-name] pbr-statistics`
Parameters	*map-name* — (Optional) Enter the name of a configured route map. A maximum of 140 characters.
Defaults	None
Command Mode	EXEC
Usage Information	None
Example	`OS10# show route-map map1 pbr-statistics`
Supported Releases	10.3.0E or later

Virtual Router Redundancy Protocol

VRRP allows you to form virtual routers from groups of physical routers on your local area network (LAN). These virtual routing platforms — master and backup pairs — provide redundancy in case of hardware failure. VRRP also allows you to easily configure a virtual router as the default gateway to all your hosts and avoids the single point of failure of a physical router.

VRRP:

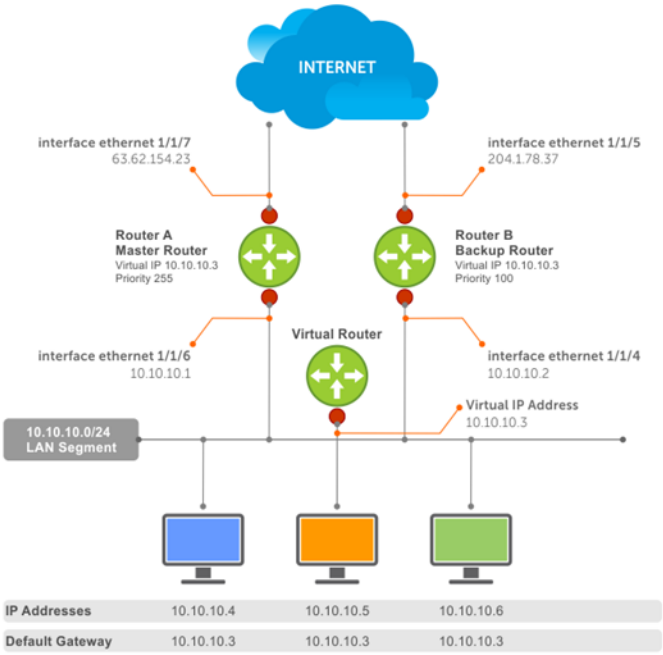
- Provides a virtual default routing platform
- Provides load balancing
- Supports multiple logical IP subnets on a single LAN segment
- Enables simple traffic routing without the single point of failure of a static default route
- Avoids issues with dynamic routing and discovery protocols
- Takes over a failed default router:
  - Within a few seconds
  - With a minimum of VRRP traffic
  - Without any interaction from hosts

Configuration

VRRP specifies a master, or active, router that owns the next-hop IP and MAC address for end stations on a LAN. The master router is chosen from the virtual routers by an election process and forwards packets sent to the next-hop IP address. If the master router fails, VRRP begins the election process to choose a new master router which continues routing traffic.

VRRP packets transmit with the virtual router MAC address as the source MAC address. The virtual router MAC address associated with a virtual router is in 00:00:5E:00:01:{VRID} format for IPv4 and 00:00:5E:00:02:{VRID} format for IPv6. The VRID is the virtual router identifier that allows up to 255 IPv4 and IPv6 VRRP routers on a network. The first four octets are unquenchable, the last two octets are 01:{VRID} for IPv4 and 02:{VRID} for IPv6. The final octet changes depending on the VRRP virtual router identifier.

Basic VRRP Configuration



The example shows a typical network configuration using VRRP. Instead of configuring the hosts on network 10.10.10.0 with the IP address of either Router A or Router B as the default router, the default router of all hosts is set to the IP address of the virtual router. When any host on the LAN segment requests Internet access, it sends packets to the IP address of the virtual router.

Router A is configured as the master router with the virtual router IP address and sends any packets addressed to the virtual router to the Internet. Router B is the backup router and is also configured with the virtual router IP address.

If Router A, the master router, becomes unavailable, Router B, the backup router, automatically becomes the master router and responds to packets sent to the virtual IP address. All workstations continue to use the IP address of the virtual router to transmit packets destined to the Internet. Router B receives and forwards packets on `interface ethernet 1/1/5`. Until Router A resumes operation, VRRP allows Router B to provide uninterrupted service to the users on the LAN segment accessing the Internet.

Create virtual router

VRRP uses the VRID to identify each virtual router configured. Before using VRRP, you must configure the interface with the primary IP address and enable it.

- Create a virtual router for the interface with the VRRP identifier in INTERFACE mode, from 1 to 255.
  ```
  vrrp-group vrrp-id
  ```
- Delete a VRRP group in INTERFACE mode.
  ```
  no vrrp-group vrrp-id
  ```

Configure VRRP

```
OS10(config)# interface ethernet 1/1/5
OS10(conf-if-eth1/1/5)# vrrp-group 254
```

Verify VRRP

```
OS10(conf-eth1/1/5-vrid-254)# do show running-configuration
...
!
interface ethernet 1/1/5
ip address 10.10.10.1/24
!
```

```
vrrp-group 254
no shutdown
...
```

Group version

Configure a VRRP version for the system. Define either VRRPv2 — `vrrp version 2` or VRRPv3 — `vrrp version 3`.

- Configure the VRRP version for IPv4 in INTERFACE mode.

  ```
  vrrp version
  ```

Configure VRRP version 3

```
OS10(config)# vrrp version 3
```

1 Set the switch with the lowest priority to `vrrp version 2`.

2 Set the switch with the highest priority to `vrrp version 3`.

3 Set all switches from `vrrp version 2` to `vrrp version 3`.

Migrate IPv4 group from VRRPv2 to VRRPv3

```
OS10_backup_switch1(config)# vrrp version 2
OS10_backup_switch2(config)# vrrp version 2
```

Set master switch to VRRPv3

```
OS10_master_switch(config)# vrrp version 3
```

Set backup switches to VRRPv3

```
OS10_backup_switch1(config)# vrrp version 3
OS10_backup_switch2(config)# vrrp version 3
```

Virtual IP addresses

Virtual routers contain virtual IP addresses configured for that VRRP group (VRID). A VRRP group does not transmit VRRP packets until you assign the virtual IP address to the VRRP group.

To activate a VRRP group on an interface, configure at least one virtual IP address for a VRRP group. The virtual IP address is the IP address of the virtual router and does not require an IP address mask. You can configure up to 10 virtual IP addresses on a single VRRP group (VRID).

These rules apply to virtual IP addresses:

- The virtual IP addresses must be in the same subnet as the primary or secondary IP addresses configured on the interface. Though a single VRRP group can contain virtual IP addresses belonging to multiple IP subnets configured on the interface, Dell EMC recommends configuring virtual IP addresses belonging to the same IP subnet for any one VRRP group. An interface on which you enable VRRP contains a primary IP address of 50.1.1.1/24 and a secondary IP address of 60.1.1.1/24. The VRRP group (VRID 1) must contain virtual addresses belonging to subnet 50.1.1.0/24 or subnet 60.1.1.0/24.
- If the virtual IP address and the interface's primary/secondary IP address are the same, the priority of the VRRP group is set to 255 by default. The interface then becomes the owner router of the VRRP group and the interface's physical MAC address changes to that of the owner VRRP group's MAC address.
- If you configure multiple VRRP groups on an interface, only one of the VRRP groups can contain the interface primary or secondary IP address.

Configure virtual IP address

Configure the virtual IP address — the primary IP address and the virtual IP addresses must be on the same subnet.

1 Configure a VRRP group in INTERFACE mode, from 1 to 255.

```
vrrp-group vrrp-id
```

2 Configure virtual IP addresses for this VRRP ID in INTERFACE-VRRP mode. A maximum of 10 IP addresses.

```
virtual-address ip-address1 [...ip-address10]
```

Configure virtual IP address

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# ip address 10.1.1.1/24
OS10(conf-if-eth1/1/1)# vrrp-group 10
OS10(conf-eth1/1/1-vrid-10)# virtual-address 10.1.1.8
```

Verify virtual IP address

```
OS10# show running-configuration
! Version 10.1.9999P.2281
! Last configuration change at Jul  26 12:01:58 2016
!
aaa authentication system:local
!
interface ethernet1/1/1
 ip address 10.1.1.1/24
 no switchport
 no shutdown
 !
 vrrp-group 10
  virtual-address 10.1.1.8
!
interface ethernet1/1/2
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/3
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/4
 switchport access vlan 1
--more--
```

View VRRP information

When the VRRP process completes initialization, the State field contains either `master` or `backup`.

```
OS10# show vrrp brief
Interface      Group   Priority  Preemption  State Master-addr  Virtual addr(s)
-------------------------------------------------------------------------------
ethernet1/1/1  IPv4 10 100       true       master 10.1.1.8    10.1.1.8
```

View VRRP group 1

```
OS10# show vrrp 1
Interface : ethernet1/1/1    IPv4 VRID : 1
Primary IP Address : 10.1.1.1    State : master-state
Virtual MAC Address : 00:00:5e:00:01:01
Version : version-3    Priority : 100
Preempt :     Hold-time :
Authentication : no-authentication
Virtual IP address :
```

```
10.1.1.1
master-transitions : 1    advertise-rcvd : 0
advertise-interval-errors : 0    ip-ttl-errors : 0
priority-zero-pkts-rcvd : 0    priority-zero-pkts-sent : 0
invalid-type-pkts-rcvd : 0    address-list-errors : 0
pkt-length-errors : 0
```

Configure virtual IP address in a VRF

You can configure a VRRP group in a non-default VRF instance and assign a virtual address to this group.
To configure VRRP under a specific VRF:

1 Create the non-default VRF in which you want to configure VRRP.

```
ip vrf vrf-name
```

CONFIGURATION Mode

2 In the VRF Configuration mode, enter the desired interface.

```
interface interface-id
```

VRF CONFIGURATION Mode

3 Remove the interface from L2 switching mode.

```
no switchport
```

INTERFACE CONFIGURATION Mode

4 Assign the interface to the non-default VRF that you have created.

```
ip vrf forwarding vrf-name
```

INTERFACE CONFIGURATION Mode

5 Assign an IP address to the interface.

```
ip address ip-address
```

INTERFACE CONFIGURATION Mode

6 Configure a VRRP group.

```
vrrp-group group-id
```

INTERFACE CONFIGURATION Mode

7 Configure virtual IP address for the VRRP ID.

```
virtual-address ip-address
```

INTERFACE VRRP Mode

```
OS10(config)# ip vrf vrf-test
OS10(config-vrf)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# ip vrf forwarding vrf-test
OS10(conf-if-eth1/1/1)# ip address 10.1.1.1/24
OS10(conf-if-eth1/1/1)# vrrp-group 10
OS10(conf-eth1/1/1-vrid-10)# virtual-address 10.1.1.8
```

Before removing an interface from a VRF, delete the configured VRRP groups from the interface associated with the VRF. If you do not delete the configured VRRP groups, these groups remain active on the default VRF resulting in duplicate virtual IP address configurations.

Set group priority

The router that has the highest primary IP address of the interface becomes the *master*. The default priority for a virtual router is 100. If the master router fails, VRRP begins the election process to choose a new master router based on the next-highest priority. The virtual router priority is automatically set to 255, if any of the configured virtual IP addresses matches the interface IP address.

1 Create a virtual router for the interface with the VRRP identifier in INTERFACE mode, from 1 to 255.

```
vrrp-group vrrp-id
```

2 Configure the priority number for the VRRP group in INTERFACE-VRRP mode, from 1 to 254, default 100.

```
priority number
```

Set VRRP group priority

```
OS10(config)# interface ethernet 1/1/5
OS10(conf-if-eth1/1/5)# vrrp-group 254
OS10(conf-eth1/1/5-vrid-254)# priority 200
```

Verify VRRP group priority

```
OS10(conf-eth1/1/5-vrid-254)# do show vrrp 254

Interface : ethernet1/1/5    IPv4 VRID : 254
Primary IP Address : 10.1.1.1    State : master-state
Virtual MAC Address : 00:00:5e:00:01:01
Version : version-3    Priority : 200
Preempt :     Hold-time :
Authentication : no-authentication
Virtual IP address :
10.1.1.1
master-transitions : 1    advertise-rcvd : 0
advertise-interval-errors : 0    ip-ttl-errors : 0
priority-zero-pkts-rcvd : 0    priority-zero-pkts-sent : 0
invalid-type-pkts-rcvd : 0    address-list-errors : 0
pkt-length-errors : 0
```

Authentication

Simple authentication of VRRP packets ensures that only trusted routers participate in VRRP processes. When you enable authentication, OS10 includes the password in its VRRP transmission. The receiving router uses that password to verify the transmission.

You must configure all virtual routers in the VRRP group with the same password. You must enable authentication with the same password or authentication is disabled. Authentication for VRRPv3 is not supported.

1 Create a virtual router for the interface with the VRRP identifier in INTERFACE mode, from 1 to 255.

```
vrrp-group vrrp-id
```

2 Configure a simple text password in INTERFACE-VRRP mode.

```
authentication-type simple–text text
```

`simple–text text` — Enter the keyword and a simple text password.

Configure VRRP authentication

```
OS10(config)# interface ethernet 1/1/5
OS10(conf-if-eth1/1/5)# vrrp-group 250
OS10(conf-eth1/1/5-vrid-250)# authentication simple-text eureka
```

Verify VRRP authentication configuration

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# vrrp-group 1
OS10(conf-eth1/1/1-vrid-1)# authentication simple-text dell
```

Disable preempt

Prevent the Backup router with the higher priority from becoming the master router by disabling the preemption process. The `preempt` command is enabled by default. The command forces the system to change the master router if another router with a higher priority comes online.

You must configure all virtual routers in the VRRP group with the same settings. Configure all routers with preempt enabled or configure all with preempt disabled.

1 Create a virtual router for the interface with the VRRP identifier in INTERFACE mode, from 1 to 255.

```
vrrp-group vrrp-id
```

2 Prevent any backup router with a higher priority from becoming the Master router in INTERFACE-VRRP mode.

```
no preempt
```

Disable preempt

```
OS10(config)# interface ethernet 1/1/5
OS10(conf-if-eth1/1/5)# vrrp-group 254
OS10(conf-eth1/1/5-vrid-254)# no preempt
```

View running configuration

```
OS10(conf-eth1/1/5-vrid-254)# do show running-configuration
! Version 10.2.0E
! Last configuration change at Sep  24
07:17:45 2016
!
debug radius false
snmp-server contact http://www.dell.com/support/softwarecontacts
snmp-server location "United States"
username admin password $6$q9QBeYjZ$jfxzVqGhkxX3smxJSH9DDz7/3OJc6m5wjF8nnLD7/VKx8SloIhp4NoGZs0I/
UNwh8WVuxwfd9q4pWIgNs5BKH.
aaa authentication system:local
!
interface ethernet1/1/5
 ip address 1.1.1.1/16
 no switchport
 no shutdown
 !
 vrrp-group 254
  priority 125
  virtual-address 1.1.1.3
  no preempt
!
```

Advertisement interval

By default, the master router transmits a VRRP advertisement to all members of the VRRP group every one second, indicating it is operational and is the master router.

If the VRRP group misses three consecutive advertisements, the election process begins and the backup virtual router with the highest priority transitions to master. To avoid throttling VRRP advertisement packets, Dell EMC recommends increasing the VRRP advertisement

interval to a value higher than the default value of one second. If you change the time interval between VRRP advertisements on one router, change it on all participating routers.

If you configure VRRP version 2, you must configure the timer values in multiple of whole seconds. For example, a timer value of 3 seconds or 300 centisecs is valid and equivalent. A time value of 50 centisecs is invalid because it not a multiple of 1 second. If you are using VRRP version 3, you must configure the timer values in multiples of 25 centisecs. A centisecs is 1/100 of a second.

- Create a virtual router for the interface with the VRRP identifier in INTERFACE mode, from 1 to 255.
  `vrrp-group vrrp-id`
- For VRRPv2, change the advertisement interval setting in seconds in INTERFACE-VRRP mode, from 1 to 255, default 1.
  `advertise-interval seconds`
- For VRRPv3, change the advertisement centisecs interval setting INTERFACE-VRRP mode, from 25 to 4075, default 100.
  `advertise-interval centisecs centisecs`

Change advertisement interval

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# vrrp-group 1
OS10(conf-eth1/1/1-vrid-1)# advertise-interval centisecs 200
```

View running configuration

```
OS10(conf-eth1/1/1-vrid-1)# do show running-configuration

! Version 10.1.9999P.2281
! Last configuration change at Jul  26 12:22:33 2016
!
aaa authentication system:local
!
interface ethernet1/1/1
 ip address 10.1.1.1/16
 no switchport
 no shutdown
 !
 vrrp-group 1
  advertisment-interval centisecs 200
  priority 200
  virtual-address 10.1.1.1
!
interface ethernet1/1/2
 switchport access vlan 1
 no shutdown
```

Interface/object tracking

You can monitor the state of any interface according to the virtual group. OS10 supports a maximum of 10 track groups and each track group can track only one interface.

If the tracked interface goes down, the VRRP group's priority decreases by a default value of 10 — also known as *cost*. If the tracked interface's state goes up, the VRRP group's priority increases by the priority cost.

The lowered priority of the VRRP group may trigger an election. As the master/backup VRRP routers are selected based on the VRRP group's priority, tracking features ensure that the best VRRP router is the master for that group. The priority cost of the tracking group must be less than the configured priority of the VRRP group. If you configure the VRRP group as the owner router with a priority 255, tracking for that group is disabled, regardless of the state of the tracked interfaces. The priority of the owner group always remains 255.

For a virtual group, track the line-protocol state of any interface using the `interface` command. Enter an interface type and `node/slot/port[:subport]` information, or VLAN number:

- `ethernet` — Physical interface, from 1 to 48

- `vlan` — VLAN interface, from 1 to 4093

For a virtual group, track the status of a configured object using the `track` command and the object number. You can also configure a tracked object for a VRRP group with this command before you create the tracked object. No changes in the VRRP group's priority occur until the tracked object is determined to be down.

Configure tracking

To track the objects in a VRRP group or on interfaces, use these commands. The sum of all the costs for all tracked interfaces must be less than the configured priority of the VRRP group.

1 Assign an object tracking unique ID number in CONFIGURATION mode, from 1 to 500.

```
track track-id
```

2 Monitor an interface in Track CONFIGURATION mode.

```
interface ethernet node/slot/port[:subport]
```

Configure interface tracking

```
OS10(config)# track 10
OS10(conf-track-10)# interface ethernet 1/1/5 line-protocol
```

View running configuration

```
OS10(conf-track-10)# do show running-configuration

! Version 10.1.9999P.2281
! Last configuration change at Jul  27 03:24:01 2016
!
aaa authentication system:local
!
interface ethernet1/1/1
 ip address 10.1.1.1/16
 no switchport
 no shutdown
 !
 vrrp-group 1
  priority 200
  virtual-address 10.1.1.1
!
interface ethernet1/1/2
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/3
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/4
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/5
 switchport access vlan 1
 no shutdown
!
interface ethernet1/1/6
 switchport access vlan 1
 no shutdown
!
.....
.....
interface vlan1
 no shutdown
!
```

```
interface mgmt1/1/1
 no shutdown
!
support-assist
!
track 10
 interface ethernet1/1/5 line-protocol
```

To associate a track object with a VRRP group, use the `track` command inside VRRP GROUP CONFIGURATION mode.

VRRP commands

advertise-interval

Sets the time interval between VRRP advertisements.

Syntax `advertise-interval [seconds | centisecs centisecs]`

Parameters

- *seconds* — Set the advertise interval in seconds, from 1 to 255.
- `centisecs` *centisecs* — (Optional) Enter a value in multiples of 25, from 25 to 4075.

Default 1 second or 100 centisecs

Command Mode INTERFACE-VRRP

Usage Information Dell EMC recommends keeping the default setting for this command. If you change the time interval between VRRP advertisements on one router, change it on all routers. The `no` version of this command sets the VRRP advertisements timer interval back to its default value, 1 second or 100 centisecs.

Example

```
OS10(conf-eth1/1/6-vrid-250)# advertise-interval 120 centisecs 100
```

Supported Releases 10.2.0E or later

authentication-type

Enables authentication of VRRP data exchanges.

Syntax `authentication-type simple-text password`

Parameters `simple-text` *password* — Enter a simple text password.

Default Disabled

Command Mode INTERFACE-VRRP

Usage Information With authentication enabled, OS10 ensures that only trusted routers participate in routing in an autonomous network. The `no` version of this command disables authentication of VRRP data exchanges.

Example

```
OS10(conf-ethernet1/1/6-vrid-250)# authentication simple-text eureka
```

Supported Releases 10.2.0E or later

preempt

Permits or preempts a backup router with a higher priority value to become the master router.

Syntax	`preempt`
Parameters	None
Default	Enabled
Command Mode	INTERFACE-VRRP
Usage Information	VRRP uses preempt to determine what happens after a VRRP backup router becomes the master. With preempt enabled by default, VRRP switches to a backup if that backup router comes online with a priority higher than the new master router. If you disable preempt, VRRP switches only if the original master recovers or the new master fails. The `no` version of this command disables preemption.
Example	`OS10(conf-eth1/1/5-vrid-254)# preempt`
Supported Releases	10.2.0E or later

priority

Assigns a VRRP priority value for the VRRP group. The VRRP uses this value during the master election process.

Syntax	`priority` *`number`*
Parameters	*`number`* — Enter a priority value, from 1 to 254.
Default	100
Command Mode	INTERFACE-VRRP
Usage Information	To guarantee that a VRRP group becomes master, configure the priority of the VRRP group to the 254, which is the highest priority or configure the VRRP group's virtual IP address with same IP address as the interface's primary IP address. If you set the priority to 254 and the `virtual-address` is not equal to the interface's primary IP address, the system displays an error message. The `no` version of this command resets the value to the default of 100.
Example	`OS10(conf-eth1/1/5-vrid-254)# priority 200`
Supported Releases	10.2.0E or later

show vrrp

Displays VRRP group information.

Syntax	`show vrrp [vrf` *`vrf-name`*`] {brief |` *`vrrp-id`* `| ipv6` *`group-id`*`}`
Parameters	• `vrf` *`vrf-name`* — Displays the VRRP group information corresponding to the specified VRF. • `brief` — Displays the configuration information for all VRRP instances in the system. • *`vrrp-id`* — Enter a VRRP group ID number to view the VRRP IPv4 group operational status information, from 1 to 255. • `ipv6` *`group-id`* — (Optional) Enter a VRRP group ID number to view the specific IPv6 group operational status information, from 1 to 255.

Default All IPv4 VRRP group configuration

Command Mode EXEC

Usage Information Displays all active VRRP groups. If no VRRP groups are active, the system displays `No Active VRRP group`.

Example (Brief)

```
OS10 # show vrrp brief
Interface     Group Priority Preemption State    Master-addr Virtual addr(s)
--------------------------------------------------------------------------
ethernet1/1/1 1     200      true      master-state 10.1.1.1  10.1.1.1
```

Example (IPv6)

```
OS10 # show vrrp ipv6 1
Interface : ethernet1/1/1    IPv6 VRID : 1
Primary IP Address : 10::1    State : master-state
Virtual MAC Address : 00:00:5e:00:02:01
Version : version-3    Priority : 200
Preempt :     Hold-time :
Authentication : no-authentication
Virtual IP address :
10::1
master-transitions : 1    advertise-rcvd : 0
advertise-interval-errors : 0    ip-ttl-errors : 0
priority-zero-pkts-rcvd : 0    priority-zero-pkts-sent : 0
invalid-type-pkts-rcvd : 0    address-list-errors : 0
pkt-length-errors : 0
```

Supported Releases 10.2.0E or later

track

Assigns a unique identifier to track an object.

Syntax `track *track-id* [priority cost [value]]`

Parameters

- *track-id* — Enter the object tracking resource ID number, from 1 to 500.
- `priority cost` *value* — (Optional) Enter a cost value to subtract from the priority value, from 1 to 254.

Default 10

Command Mode INTERFACE-VRRP

Usage Information If you disable the interface, the cost value subtracts from the priority value and forces a new master election. This election process is applicable when the priority value is lower than the priority value in the backup virtual router. The `no` version of this command resets the value to the default.

Example

```
OS10(conf-eth1/1/5-vrid-254)# track 400
```

Example (Priority Cost)

```
OS10(conf-eth1/1/5-vrid-254)# track 400 priority-cost 20
```

Supported Releases 10.2.0E or later

track interface

Monitors an interface and lowers the priority value of the VRRP group on that interface, if disabled.

Syntax `interface {ethernet *node/slot/port[:subport]*} [line-protocol]`

Parameters

- `ethernet` *node/slot/port[:subport]* — (Optional) Enter the keyword and the interface information to track.
- `line-protocol` — (Optional) Tracks the interface line-protocol operational status.

Default Disabled

Command Mode EXEC

Usage Information Assign an object tracking unique ID number before tracking the interface. Use the `line-protocol` parameter to track for interface operational status information. The `no` version of this command resets the value to the default.

Example

```
OS10(config)# track 10
OS10(conf-track-10)# interface ethernet 1/1/5 line-protocol
```

Supported Releases 10.2.0E or later

virtual-address

Configures up to 10 virtual router IP addresses in the VRRP group. Set at least one virtual IP address for the VRRP group to start sending VRRP packets.

Syntax `virtual-address` *ip-address1 [ip-address2...ip-address10]*

Parameters

- *ip-address1* — Enter the IP address of a virtual router in A.B.C.D format. The IP address must be on the same subnet as the interface's primary IP address.
- *ip-address2...ip-address10* — (Optional) Enter up to nine additional IP addresses of virtual routers, separated by a space. The IP addresses must be on the same subnet as the interface's primary IP address.

Default Enabled

Command Mode INTERFACE-VRRP

Usage Information The VRRP group only becomes active and sends VRRP packets when you configure a virtual IP address. When you delete the virtual address, the VRRP group stops sending VRRP packets. To guarantee that a VRRP group becomes master, configure the VRRP group's virtual address with the same IP address as the interface's primary IP address and change the priority of the VRRP group to 255. You can ping the virtual addresses configured in all VRRP groups. The `no` version of this command deletes one or more virtual-addresses configured in the system.

Example

```
OS10(conf-eth1/1/5-vrid-254)# virtual address 10.1.1.15
```

Supported Releases 10.2.0E or later

vrrp delay reload

Sets the delay time for VRRP initialization after a system reboot.

Syntax `vrrp delay reload` *seconds*

Parameters *seconds* — Enter the number of seconds for the VRRP reload time, from 0 to 900.

Default 0

Command Mode CONFIGURATION

Usage Information VRRP delay reload time of zero seconds indicates no delays. This command configuration applies to all the VRRP configured interfaces. The `no` version of this command resets the value to the default.

Example	`OS10(config)# vrrp delay reload 5`
Supported Releases	10.4.0E(R1) or later

vrrp-group

Assigns a VRRP group identification number to an IPv4 interface or VLAN

Syntax	`vrrp-group` *vrrp-id*
Parameters	*vrrp-id* — Enter a VRRP group identification number, from 1 to 255.
Default	Not configured
Command Mode	INTERFACE-VRRP
Usage Information	The VRRP group only becomes active and sends VRRP packets when you configure a virtual IP address. When you delete the virtual address, the VRRP group stops sending VRRP packets. The `no` version of this command removes the `vrrp-group` configuration.
Example	`OS10(conf-if-eth1/1/5)# vrrp-group 254`
Example (VLAN)	`OS10(conf-if-vl-10)# vrrp-group 5`
Supported Releases	10.2.0E or later

vrrp-ipv6-group

Assigns a VRRP group identification number to an IPv6 interface.

Syntax	`vrrp-ipv6-group` *vrrp-id*
Parameters	*vrrp-id* — Enter a VRRP group identification number, from 1 to 255.
Default	Not configured
Command Mode	INTERFACE-VRRP
Usage Information	The VRRP group only becomes active and sends VRRP packets when you configure a virtual IP address. When you delete the virtual address, the VRRP group stops sending VRRP packets. The `no` version of this command removes the `vrrp-ipv6-group` configuration.
Example	`OS10(conf-if-eth1/1/7)# vrrp-ipv6-group 250`
Supported Releases	10.2.0E or later

vrrp version

Sets the VRRP version for the IPv4 group.

Syntax	`vrrp version {2 | 3}`
Parameters	• `2` — Set to VRRP version 2. • `3` — Set to VRRP version 3.

Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command disables the VRRP version for the IPv4 group.
Example	`OS10(config)# vrrp version 2`
Supported Releases	10.2.0E or later

6

Multicast

Multicast is a technique that allows networking devices to send data to a group of interested receivers in a single transmission. For instance, this technique is widely used for streaming videos. Multicast allows you to more efficiently use network resources, specifically for bandwidth-consuming services such as audio and video transmission.

OS10 supports the multicast feature in IPv4 networks and uses the following protocols for multicast distribution:

- Internet Group Management Protocol (IGMP)
- Protocol Independent Multicast (PIM)

Important notes

- OS10 supports IGMP and IPv4 PIM for multicast routing. This release of OS10 does not support IPv6 PIM.
- OS10 supports PIM and IGMP on default and non-default VRFs.
- OS10 does not support multicast routing on S5148F-ON and S3048-ON platforms.
- Multicast flood control is not supported on S4248FB-ON, S4248FBL-ON, and S5148F-ON platforms.

Configure multicast

Configuring multicast is a two-step process that involves configuring multicast routing and enabling PIM sparse mode (PIM-SM) on a Layer 3 (L3) interface. The following procedure describes how to configure multicast routing.

For more information about IGMP and PIM feature configurations, see Internet Group Management Protocol and Protocol Independent Multicast.

NOTE:
Multicast flood restrict feature is enabled by default. To ensure that no traffic drops occur, Dell EMC recommends that you do one of the following:

- Disable IGMP snooping on the VLAN between two PIM routers that do not have IGMP receivers on that VLAN.
- Configure the interface between the PIM routers as static mrouter port.

1 Enable multicast routing for IPv4 networks.

```
OS10# configure terminal
OS10(config)# ip multicast-routing
```

2 Configure an IP address to a VLAN interface.

```
OS10(config)# interface vlan 2
OS10(conf-if-vl-2)# ip address 1.1.1.2/24
```

3 Enable PIM sparse mode on an L3 interface.

```
OS10(config)# interface vlan 2
OS10(conf-if-vl-2) ip pim sparse-mode
```

4 From CONFIGURATION mode, configure the rendezvous point (RP) IP address statically and specify the multicast group address range. The RP IP address should be reachable across the PIM domain.

```
OS10(config)# ip pim rp-address 171.1.1.1 group-address 225.1.1.3/32
```

Configure the RP address and multicast group address on all nodes in your network.

Unknown multicast flood control

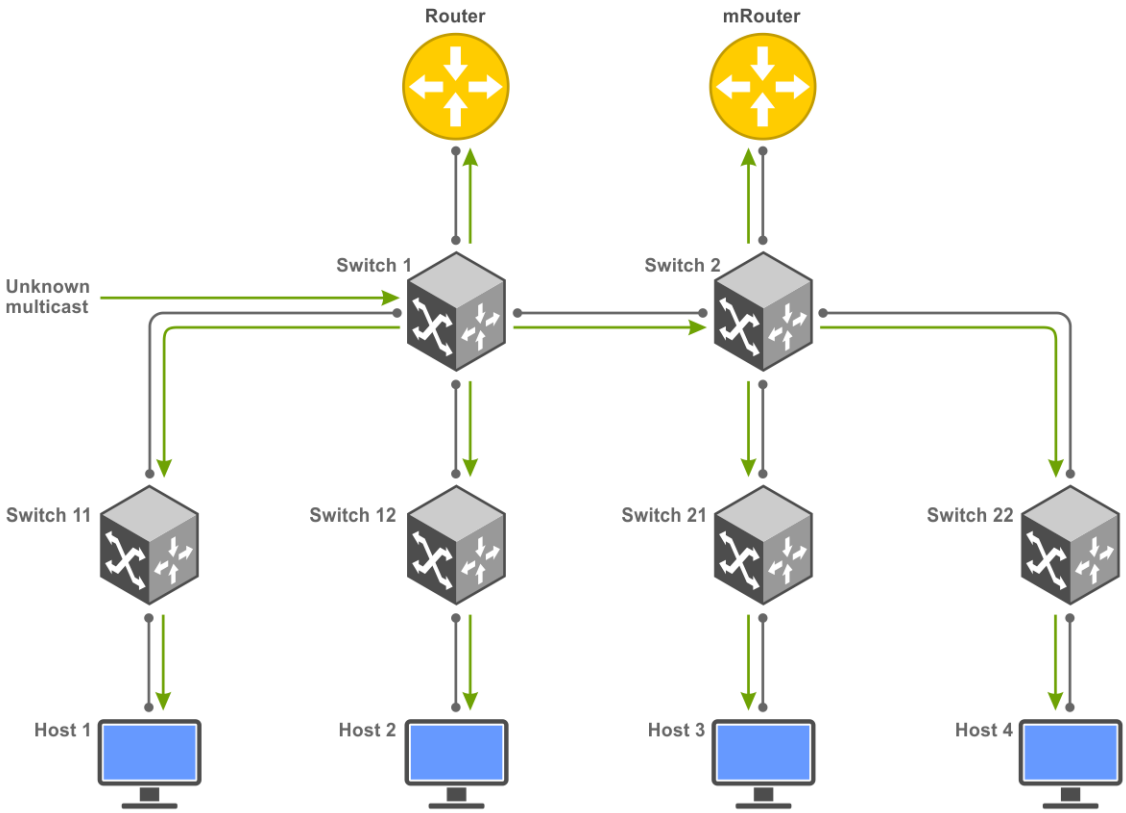
The unknown multicast flood control feature enables the system to forward unknown multicast packets only to a multicast router (mrouter).

When you enable multicast snooping, OS10 forwards multicast frames, whose destination is already learned, to their intended recipients. When the system receives multicast frames whose destination is not known, it floods the frames for all ports on the specific VLAN. All hosts that receive these multicast frames must process them. With multicast flood control, the system forwards unknown multicast frames only to the interface that leads to the mrouter. The mrouter can then forward the traffic to the intended destinations.

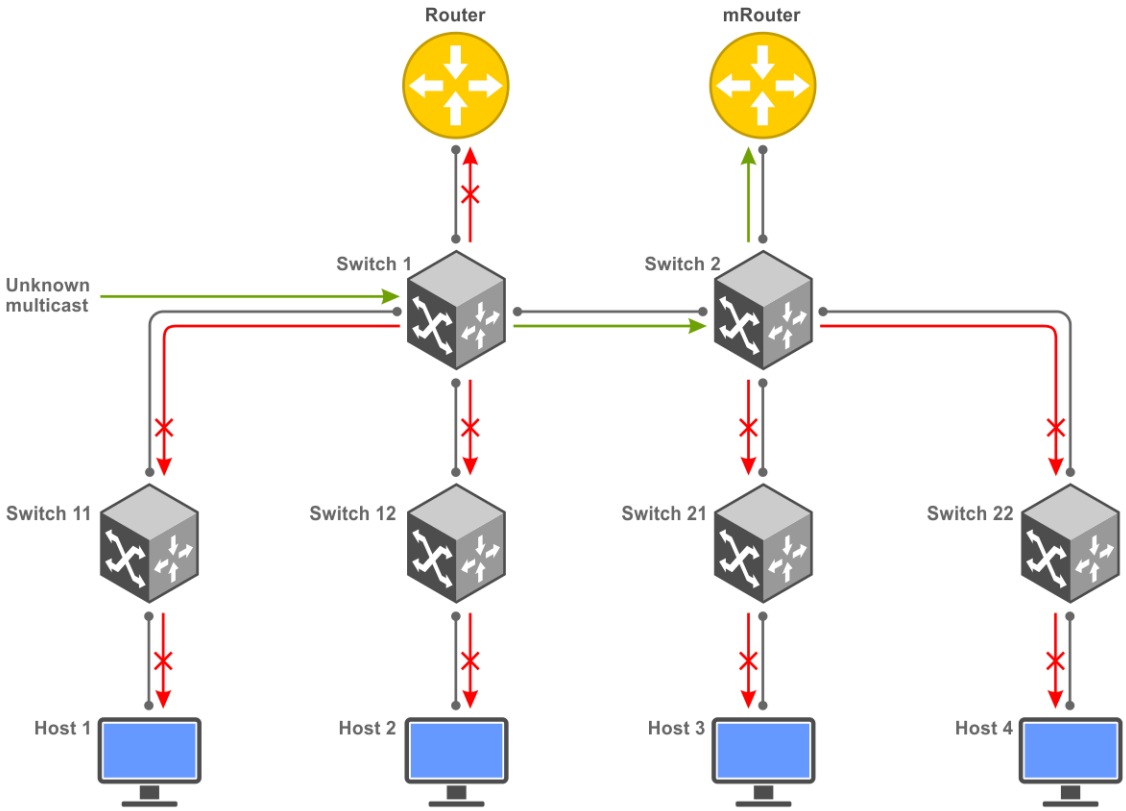
For multicast flood control to work, you must enable both IGMP and MLD snooping on the system. By default, multicast flood control, IGMP snooping, and MLD snooping are enabled.

NOTE: The Multicast flood control feature is not supported on the Dell EMC Networking S4248FB-ON, S4248FBL-ON, and S5148F-ON switches.

The following describes a scenario where a multicast frame is flooded on all ports of all switches. The switches and hosts in the network need not receive these frames because they are not the intended destinations.



With multicast flood control, multicast frames, whose destination is not known, are forwarded only to the designated mrouter port. OS10 learns of the mrouter interface dynamically based on the interface where an IGMP membership query is received. You can also statically configure the mrouter interface using the `ip igmp snooping mrouter` and `ipv6 mld snooping mrouter` commands.



Enable multicast flood control

Multicast flood control is enabled on OS10 by default. If it is disabled, use the following procedure to enable multicast flood control:

1 Configure IGMP snooping. To know how to configure IGMP snooping, see the IGMP snooping section.
2 Configure MLD snooping. To know how to configure MLD snooping, see the MLD Snooping section.
3 Enable the multicast flood control feature.

```
OS10(config)# multicast snooping flood-restrict
```

Multicast Commands

multicast snooping flood-restrict

Enables multicast snooping flood control for IGMP snooping and MLD snooping.

Syntax	`multicast snooping flood-restrict` The `no` version of this command disables multicast flood control.
Parameters	None
Default	Enabled
Command Mode	CONFIGURATION

Usage Information	Multicast snooping flood control, IGMP snooping, and MLD snooping are enabled by default. For multicast flood restrict to be effective on a VLAN, IGMP snooping and MLD snooping must be enabled at both global and VLAN levels. To disable multicast snooping flood control, use the `no multicast snooping flood-restrict` command.
Example	`OS10(config)# multicast snooping flood-restrict`
Supported Releases	10.4.3.0 or later

Internet Group Management Protocol

Internet Group Management Protocol (IGMP) is a communications protocol that establishes multicast group memberships using IPv4 networks. OS10 supports IGMPv1, IGMPv2, and IGMPv3 to manage the multicast group memberships on IPv4 networks.

The IGMP querier periodically (by default, every 60 seconds) sends out a membership query to all the hosts. The hosts, in response to the query, send a response back to the querier to report their multicast group memberships. The switch makes an entry to identify the corresponding port as a member of the particular multicast group..

NOTE: A multicast router is a Layer 3 router or switch that has multicast features enabled.

When a host wants to join a multicast group, it sends an IGMP message to the multicast router.

Each network segment has an IGMP querier, which is a multicast router. The multicast router periodically sends IGMP queries to learn which multicast groups are active and have members on the network.

Multicast routers send the following types of queries:

- General query—To learn about listeners for multicast groups.
- Multicast address-specific query—To learn if a particular multicast address has listeners.
- Multicast address-and-source-specific query—To learn if any of the sources from the specified list for a multicast source has any listeners.

The hosts send the following messages to multicast routers:

- Version 1: Membership report
- Version 2:
  - Version 1 membership report for backward compatibility with version 1
  - Version 2 membership report
  - Leave group message
- Version 3:
  - Version 1 membership report for backward compatibility with version 1
  - Version 2 membership report for backward compatibility with version 2
  - Version 3 membership report
  - Version 2 leave group message

  Version 3 provides support for source filtering. The system reports interest in receiving packets only from specific source addresses, or from all the sources except some specific source addresses, sent to a particular multicast address.

Standards compliance

- OS10 complies to the RFCs 1112, 2236, and 3376 for IGMP versions 1, 2, and 3, respectively.

- OS10 uses version 3 as the default IGMP version. Version 3 is backwards compatible with versions 1 and 2.

Important notes

- OS10 systems cannot serve as an IGMP host or an IGMP version 1 querier.
- OS10 automatically enables IGMP on interfaces where you enable PIM sparse mode.

Supported IGMP versions

IGMP has three versions. Version 3 obsoletes and is backwards-compatible with version 2; version 2 obsoletes version 1.
OS10 supports the following IGMP versions:

- Router—IGMP versions 2 and 3. The default is version 3.
- Host—IGMP versions 1, 2, and 3.

In IGMP version 2, the host expresses interest in a particular group membership (*, G). In IGMP version 3, the host expresses interest in a particular group membership, and specifies the source from which it wants the multicast traffic (S, G).

Query interval

The IGMP querier periodically sends a general query to discover which multicast groups are active. A group must have at least one host to be active. By default, the periodic query messages are sent every 60 seconds. You can configure this value using the `ip igmp query-interval` command.

To configure a query interval:

```
OS10# configure terminal
OS10# interface vlan120
OS10(conf-if-vl-120)# ip igmp query-interval 60
```

Last member query interval

When the IGMP querier receives a leave message, it sends a group-specific query message to ensure if any other host in the network is interested in the multicast flow. By default, the group-specific query messages are sent every 1000 milliseconds. You can configure this value using the `ip igmp last-member-query-interval` command.

To configure last member query interval:

```
OS10# configure terminal
OS10# interface vlan120
OS10(conf-if-vl-120)# ip igmp last-member-query-interval 200
```

Maximum response time

The maximum response time is the amount of time that the querier waits for a response to a query before taking action.

When a host receives a query, it does not respond immediately, but rather starts a delay timer. The delay time is set to a random value between 0 and the maximum response time. The host sends a response when the timer expires; in IGMP version 2, if another host responds before the timer expires, the timer nullifies, and no response is sent.

The querier advertises the maximum response time in the query. Lowering this value decreases leave latency but increases response burstiness because all host membership reports are sent before the maximum response time expires. Inversely, increasing this value decreases burstiness, but increases leave latency.

To configure maximum response time:

```
OS10# configure terminal
OS10# interface vlan120
OS10(conf-if-vl-120)# ip igmp query-max-resp-time 20
```

IGMP immediate leave

If the IGMP querier does not receive a response to a group-specific or group-and-source query, it sends another query based on the configured querier robustness value. This value determines the number of times the querier sends the message. If the querier does not receive a response, it removes the group from the outgoing interface for the subnet.

IGMP immediate leave reduces leave latency by enabling a router to immediately delete the group membership on an interface after receiving a *leave* message. Immediate leave does not send group-specific or group-and-source queries before deleting the entry.

To configure IGMP immediate leave:

```
OS10# configure terminal
OS10# interface vlan1
OS10(conf-if-vl-1)# ip igmp immediate-leave
```

Select an IGMP version

OS10 enables IGMP version 3 by default.

If hosts require an IGMP version other than 3, use the following to select a different IGMP version:

```
OS10# configure terminal
OS10# interface vlan1
OS10(conf-if-vl-1)# ip igmp version 3
```

View IGMP-enabled interfaces and groups

To view IGMP-enabled interfaces and groups, use the following `show` commands.

To view IGMP-enabled interfaces:

```
OS10# show ip igmp interface
Vlan103 is up, line protocol is up
Internet address is 2.1.1.2
IGMP is enabled on interface
IGMP version is 3
IGMP query interval is 60 seconds
IGMP querier timeout is 130 seconds
IGMP last member query response interval is 1000 ms
```

```
IGMP max response time is 10 seconds
IGMP immediate-leave is disabled on this interface
IGMP joins count: 0
IGMP querying router is 2.1.1.1

Vlan105 is up, line protocol is up
Internet address is 3.1.1.2
IGMP is enabled on interface
IGMP version is 3
IGMP query interval is 60 seconds
IGMP querier timeout is 130 seconds
IGMP last member query response interval is 1000 ms
IGMP max response time is 10 seconds
IGMP immediate-leave is disabled on this interface
IGMP joins count: 0
IGMP querying router is 3.1.1.1

Vlan121 is up, line protocol is up
Internet address is 121.1.1.2
IGMP is enabled on interface
IGMP version is 3
IGMP query interval is 60 seconds
IGMP querier timeout is 130 seconds
IGMP last member query response interval is 1000 ms
IGMP max response time is 10 seconds
IGMP immediate-leave is disabled on this interface
IGMP joins count: 100
IGMP querying router is 121.1.1.2
```

To view IGMP groups:

```
OS10# show ip igmp groups
Total Number of Groups: 100
IGMP Connected Group Membership
Group Address Interface     Mode           Uptime     Expires  Last Reporter
225.1.1.1     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.2     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.3     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.4     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.5     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.6     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.7     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.8     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.9     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.10    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.11    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.12    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.13    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.14    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.15    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.16    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
```

IGMP snooping

IGMP snooping uses the information in IGMP packets to generate a forwarding table that associates ports with multicast groups. When switches receive multicast frames, they forward them to their intended receivers. OS10 supports IGMP snooping on virtual local area network (VLAN) interfaces.

Effective with OS10 release 10.4.3.0, IGMP snooping is enabled by default.

Configure IGMP snooping

- Enable IGMP snooping globally using the `ip igmp snooping enable` command in CONFIGURATION mode. This command enables IGMP snooping on all VLAN interfaces.
- (Optional) Disable IGMP snooping on specific VLAN interfaces using the `no ip igmp snooping` command in VLAN INTERFACE mode.

- (Optional) Multicast flood control is enabled by default. To disable the multicast flood restrict feature, use the `no ip igmp snooping enable` command in CONFIGURATION mode. To reenable the feature globally, use the `multicast snooping flood-restrict` command in CONFIGURATION mode.
- In a network, the snooping switch is connected to a multicast Router that sends IGMP queries. On a Layer 2 network that does not have a multicast router, you can configure the snooping switch to act as querier. Use the `ip igmp snooping querier` command in VLAN INTERFACE mode to send the queries.
- OS10 learns the multicast router interface dynamically based on the interface on which IGMP membership query is received. To assign a multicast router interface statically, use the `ip igmp snooping mrouter interface` *`interface—type`* command in VLAN INTERFACE mode.

  NOTE:
  IGMP snooping dynamically detects the mrouter interface based on IGMP queries that it receives. If there are more than one multicast routers connected to the snooping switch, one of them will send IGMP queries and the interface connected to that router is dynamically learnt as an mrouter port. You must configure the interfaces connected to other multicast routers as static mrouter port.

- (Optional) Configure the IGMP version using the `ip igmp version` *`version-number`* command in VLAN INTERFACE mode.
- (Optional) The fast leave option allows the IGMP snooping switch to remove an interface from the multicast group immediately on receiving the *leave* message. Enable fast leave with the `ip igmp snooping fast-leave` command in VLAN INTERFACE mode.
- (Optional) Configure the time interval for sending IGMP general queries with the `ip igmp snooping query-interval` *`query-interval-time`* command in VLAN INTERFACE mode.
- (Optional) Configure the maximum time for responding to a query advertised in IGMP queries using the `ip igmp snooping query-max-resp-time` *`query-response-time`* command in VLAN INTERFACE mode.
- (Optional) Configures the time interval between group-specific IGMP query messages with the `ip igmp snooping last-member-query-interval` *`query-interval-time`* command in VLAN INTERFACE mode.

IGMP snooping configuration

```
OS10(config)# ip igmp snooping enable
OS10(config)# interface vlan 100
OS10(conf-if-vl-100)# ip igmp snooping mrouter interface ethernet 1/1/32
OS10(conf-if-vl-100)# ip igmp snooping querier
OS10(conf-if-vl-100)# ip igmp version 3
OS10(conf-if-vl-100)# ip igmp snooping fast-leave
OS10(conf-if-vl-100)# ip igmp snooping query-interval 60
OS10(conf-if-vl-100)# ip igmp snooping query-max-resp-time 10
OS10(conf-if-vl-100)# ip igmp snooping last-member-query-interval 1000
```

View IGMP snooping information

```
OS10# show ip igmp snooping groups
Total Number of Groups: 480
IGMP Connected Group Membership
Group Address                  Interface           Mode            Expires
225.1.0.0                      vlan3531            IGMPv2-Compat   00:01:35
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
225.1.0.1                      vlan3531            IGMPv2-Compat   00:01:35
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
225.1.0.2                      vlan3531            IGMPv2-Compat   00:01:35
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
225.1.0.3                      vlan3531            IGMPv2-Compat   00:01:35
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
225.1.0.4                      vlan3531            IGMPv2-Compat   00:01:35
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
225.1.0.5                      vlan3531            IGMPv2-Compat   00:01:35
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
225.1.0.6                      vlan3531            IGMPv2-Compat   00:01:35
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
225.1.0.7                      vlan3531            IGMPv2-Compat   00:01:35
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
225.1.0.8                      vlan3531            IGMPv2-Compat   00:01:35
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
225.1.0.9                      vlan3531            IGMPv2-Compat   00:01:35
```

```
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
--more--

<<Output Truncated>>
```

```
OS10# show ip igmp snooping interface vlan 2
Vlan2 is up, line protocol is up
IGMP version is 3
IGMP snooping is enabled on interface
IGMP snooping query interval is 60 seconds
IGMP snooping querier timeout is 130 seconds
IGMP snooping last member query response interval is 1000 ms
IGMP Snooping max response time is 10 seconds
IGMP snooping fast-leave is disabled on this interface
IGMP snooping querier is disabled on this interface
Multicast flood-restrict is enabled on this interface
```

```
show ip igmp snooping mrouter
Interface    Router Ports
Vlan 100      ethernet 1/1/32
```

IGMP commands

clear ip igmp groups

Clears entries from the group cache table.

Syntax	`clear ip igmp [vrf vrf-name] groups`
Parameters	`vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.
Default	None
Command Mode	EXEC
Usage Information	None
Example	`OS10# clear ip igmp groups`
Supported Releases	10.4.3.0 or later

ip igmp immediate-leave

Enables IGMP immediate leave.

Syntax	`ip igmp immediate-leave`
Parameters	None
Default	None
Command Mode	INTERFACE
Usage Information	The querier sends some group-specific queries when it receives a leave message before deleting the group from the membership database. If you need to immediately delete a group from the membership database, use the `ip igmp immediate-leave` command. The `no` version of this command disables IGMP immediate leave.

Example

```
OS10# configure terminal
OS10# interface vlan11
OS10(conf-if-vl-11)# ip igmp immediate-leave
```

Supported Releases 10.4.3.0 or later

ip igmp last-member-query-interval

Changes the last member query interval, which is the maximum response time included in the group-specific queries sent in response to leave group messages. This last-member-query-interval is the interval between group-specific query messages.

Syntax `ip igmp last-member-query-interval milliseconds`

Parameters *milliseconds*—Enter the amount of time in milliseconds to configure the time interval between group-specific query messages. The range is from 100 to 65535.

Default 1000 milliseconds

Command Mode INTERFACE

Usage Information None

Example

```
OS10# configure terminal
OS10# interface vlan11
OS10(conf-if-vl-11)# ip igmp last-member-query-interval 200
```

Supported Releases 10.4.3.0 or later

ip igmp query-interval

Changes the frequency of IGMP general queries sent by the querier.

Syntax `ip igmp query-interval seconds`

Parameters *seconds*—Enter the amount of time in seconds to configure the time interval for IGMP general queries. The range is from 1 to 18000.

Default 60 seconds

Command Mode INTERFACE

Usage Information None

Example

```
OS10# configure terminal
OS10# interface vlan12
OS10(conf-if-vl-12)# ip igmp query-interval 60
```

Supported Releases 10.4.3.0 or later

ip igmp query-max-resp-time

Configures the maximum query response time advertised in general queries.

Syntax `ip igmp query-max-resp-time seconds`

Parameters *seconds*—Enter the amount of time in seconds, from 1 to 25.

Default 10 seconds

Command Mode	INTERFACE
Usage Information	The IGMP query maximum response time value must be less than the IGMP query interval value. The `no` form of the command configures the default value.

Example	`OS10# configure terminal`
	`OS10# interface vlan14`
	`OS10(conf-if-vl-14)# ip igmp query-max-resp-time 20`

Supported Releases	10.4.3.0 or later

ip igmp snooping enable

Enables IGMP snooping globally.

Syntax	`ip igmp snooping enable`
Parameters	None
Default	Enabled
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command disables IGMP snooping.
Example	`OS10(config)# ip igmp snooping enable`
Supported Releases	10.4.0E(R1) or later

ip igmp snooping

Enables IGMP snooping on the specified VLAN interface.

Syntax	`ip igmp snooping`
Parameters	None
Default	Depends on the global configuration.
Command Mode	VLAN INTERFACE
Usage Information	When you enable IGMP snooping globally, the configuration applies to all VLAN interfaces. You can disable IGMP snooping on specified VLAN interfaces. The `no` version of this command disables IGMP snooping on the specified VLAN interface.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# no ip igmp snooping`
Supported Releases	10.4.0E(R1) or later

ip igmp snooping fast-leave

Enables fast leave in IGMP snooping for specified VLAN.

Syntax	`ip igmp snooping fast-leave`
Parameters	None

Default	Disabled
Command Mode	VLAN INTERFACE
Usage Information	The fast leave option allows the IGMP snooping switch to remove an interface from the multicast group immediately on receiving the *leave* message. The `no` version of this command disables the fast leave functionality.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# ip igmp snooping fast-leave`
Supported Releases	10.4.1.0 or later

ip igmp snooping last-member-query-interval

Configures the time interval between group-specific IGMP query messages.

Syntax	`ip igmp snooping last-member-query-interval` *`query-interval-time`*
Parameters	*`query-interval-time`*—Enter the query time interval in milliseconds, from 100 to 65535.
Default	1000 milliseconds
Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command resets the last member query interval time to the default value.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# ip igmp snooping last-member-query-interval 2500`
Supported Releases	10.4.1.0 or later

ip igmp snooping mrouter

Configures multicast router port on the specified VLAN interface.

Syntax	`ip igmp snooping mrouter interface` *`interface-type`*
Parameters	*`interface-type`*—Enter the interface type details. The interface must be a member of the VLAN.
Default	Not configured
Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command removes the multicast router configuration from the VLAN member port.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# ip igmp snooping mrouter interface ethernet 1/1/1`
Supported Releases	10.4.0E(R1) or later

ip igmp snooping querier

Enables IGMP querier processing for the specified VLAN interface.

Syntax	`ip igmp snooping querier`
Parameters	None
Default	Not configured

Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command disables IGMP querier on the VLAN interface..
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# ip igmp snooping querier`
Supported Releases	10.4.0E(R1) or later

ip igmp snooping query-interval

Configures time interval for sending IGMP general queries.

Syntax	`ip igmp snooping query-interval` *`query-interval-time`*
Parameters	*`query-interval-time`*—Enter the interval time in seconds, from 2 to 18000.
Default	60 seconds
Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command resets the query interval to the default value.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# ip igmp snooping query-interval 120`
Supported Releases	10.4.1.0 or later

ip igmp query-max-resp-time

Configures the maximum time for responding to a query advertised in IGMP queries.

Syntax	`ip igmp snooping query-max-resp-time` *`query-response-time`*
Parameters	*`query-response-time`*—Enter the query response time in seconds, ranging from 1 to 25.
Default	10 seconds
Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command resets the query response time to default value.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# ip igmp snooping query-max-resp-time 15`
Supported Releases	10.4.1.0 or later

ip igmp version

Configures IGMP version.

Syntax	`ip igmp version` *`version-number`*
Parameters	*`version-number`*—Enter the version number as 2 or 3.
Default	3
Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command resets the version number to the default value.

Example

```
OS10(config)# interface vlan 100
OS10(conf-if-vl-100)# ip igmp version 2
```

Supported Releases 10.4.1.0 or later

show ip igmp groups

Displays the IGMP groups.

Syntax `show ip igmp [vrf vrf-name] groups [group-address [detail] | detail | interface-name [group-address [detail]]]`

Parameters

- `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.
- *group-address*—Enter the group address in dotted decimal format to view specific group information.
- *interface-name*—Enter the interface name.

Default None

Command Mode EXEC

Usage Information The `show ip igmp groups` command displays the IGMP database, configured entries for all groups on all interfaces, all groups on specific interfaces, or specific groups on specific interfaces. This command displays the following:

- `Group address`—Lists the multicast address for the IGMP group
- `Interface`—Lists the interface type, slot, and port number
- `Mode`—Displays the IGMP version used
- `Uptime`—Displays the amount of time the group has been operational.
- `Expires`—Displays the amount of time until the entry expires
- `Last reporter`—Displays the IP address of the last host to be a member of the IGMP group

Example

```
OS10# show ip igmp groups
Total Number of Groups: 100
IGMP Connected Group Membership
Group Address Interface     Mode           Uptime     Expires  Last Reporter
225.1.1.1     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.2     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.3     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.4     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.5     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.6     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.7     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.8     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.9     vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.10    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.11    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.12    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.13    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.14    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.15    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
225.1.1.16    vlan121    IGMPv2-Compat   12:39:00    00:01:58  121.1.1.10
```

Supported Releases 10.4.3.0 or later

show ip igmp interface

Displays information about all IGMP-enabled interfaces.

Syntax `show ip igmp [vrf vrf-name] interface name`

Parameters

- `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.
- `interface name`—Enter the keyword `interface`, then the interface name.

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip igmp interface
Vlan103 is up, line protocol is up
Internet address is 2.1.1.2
IGMP is enabled on interface
IGMP version is 3
IGMP query interval is 60 seconds
IGMP querier timeout is 130 seconds
IGMP last member query response interval is 1000 ms
IGMP max response time is 10 seconds
IGMP immediate-leave is disabled on this interface
IGMP joins count: 0
IGMP querying router is 2.1.1.1

Vlan121 is up, line protocol is up
Internet address is 121.1.1.2
IGMP is enabled on interface
IGMP version is 3
IGMP query interval is 60 seconds
IGMP querier timeout is 130 seconds
IGMP last member query response interval is 1000 ms
IGMP max response time is 10 seconds
IGMP immediate-leave is disabled on this interface
IGMP joins count: 100
IGMP querying router is 121.1.1.2
```

Supported Releases 10.4.3.0 or later

show ip igmp snooping groups

Displays IGMP snooping group membership details.

Syntax `show ip igmp snooping groups [detail | [vlan vlan-id [detail | ip-address]]]`

Parameters

- `vlan-id`—(Optional) Enter the VLAN ID, from 1 to 4093.
- `detail`—(Optional) Enter `detail` to display the IGMPv3 source information.
- `ip-address`—(Optional) Enter the IP address of the multicast group.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip igmp snooping groups
Total Number of Groups: 480
IGMP Connected Group Membership
Group Address                  Interface           Mode               Expires
225.1.0.0                      vlan3031            IGMPv2-Compat      00:01:26
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.1                      vlan3031            IGMPv2-Compat      00:01:26
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.2                      vlan3031            IGMPv2-Compat      00:01:26
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.3                      vlan3031            IGMPv2-Compat      00:01:26
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.4                      vlan3031            IGMPv2-Compat      00:01:26
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.5                      vlan3031            IGMPv2-Compat      00:01:26
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.6                      vlan3031            IGMPv2-Compat      00:01:26
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.7                      vlan3031            IGMPv2-Compat      00:01:26
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.8                      vlan3031            IGMPv2-Compat      00:01:26
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.9                      vlan3031            IGMPv2-Compat      00:01:26
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.10                     vlan3031            IGMPv2-Compat      00:01:26
--more--
<<Output Truncated>>
```

Example (with VLAN)

```
OS10# show ip igmp snooping groups vlan 3031
Total Number of Groups: 12
IGMP Connected Group Membership
Group Address                  Interface           Mode                Expires
225.1.0.0                      vlan3031            IGMPv2-Compat       00:01:30
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.1                      vlan3031            IGMPv2-Compat       00:01:30
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.2                      vlan3031            IGMPv2-Compat       00:01:30
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.3                      vlan3031            IGMPv2-Compat       00:01:30
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.4                      vlan3031            IGMPv2-Compat       00:01:30
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.5                      vlan3031            IGMPv2-Compat       00:01:30
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.6                      vlan3031            IGMPv2-Compat       00:01:30
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.7                      vlan3031            IGMPv2-Compat       00:01:30
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.8                      vlan3031            IGMPv2-Compat       00:01:30
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.9                      vlan3031            IGMPv2-Compat       00:01:30
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
225.1.0.10                     vlan3031            IGMPv2-Compat       00:01:30
--more--
```

Example (with VLAN and multicast IP address)

```
OS10# show ip igmp snooping groups vlan 3031 225.1.0.0
IGMP Connected Group Membership
Group Address                  Interface           Mode                Expires
225.1.0.0                      vlan3031            IGMPv2-Compat       00:01:44
  Member-ports :port-channel51,ethernet1/1/51:1,ethernet1/1/52:1
```

Example (with detail)

```
OS10# show ip igmp snooping groups detail
Interface      vlan3041
Group          232.11.0.0
Source List
 101.41.0.21
```

```
   Member Port         Mode                Uptime              Expires
   port-channel51      Include             1d:20:26:07         00:01:41
   ethernet1/1/51:1    Include             1d:20:26:05         00:01:46
   ethernet1/1/52:1    Include             1d:20:26:08         00:01:46

Interface      vlan3041
Group          232.11.0.1
Source List
 101.41.0.21
   Member Port         Mode                Uptime              Expires
   port-channel51      Include             1d:20:26:07         00:01:41
   ethernet1/1/51:1    Include             1d:20:26:05         00:01:46
   ethernet1/1/52:1    Include             1d:20:26:08         00:01:46

Interface      vlan3041
Group          232.11.0.2
Source List
 101.41.0.21
   Member Port         Mode                Uptime              Expires
   port-channel51      Include             1d:20:26:07         00:01:41
--more-- <<Output Truncated>>
```

Example (with VLAN)

```
OS10# show ip igmp snooping groups vlan 3041 detail
Interface      vlan3041
Group          232.11.0.0
Source List
 101.41.0.21
   Member Port         Mode                Uptime              Expires
   port-channel51      Include             1d:20:26:07         00:01:41
   ethernet1/1/51:1    Include             1d:20:26:05         00:01:46
   ethernet1/1/52:1    Include             1d:20:26:08         00:01:46

Interface      vlan3041
Group          232.11.0.1
Source List
 101.41.0.21
   Member Port         Mode                Uptime              Expires
   port-channel51      Include             1d:20:26:07         00:01:41
   ethernet1/1/51:1    Include             1d:20:26:05         00:01:46
   ethernet1/1/52:1    Include             1d:20:26:08         00:01:46

Interface      vlan3041
Group          232.11.0.2
Source List
 101.41.0.21
   Member Port         Mode                Uptime              Expires
   port-channel51      Include             1d:20:26:07         00:01:41
--more--
```

Example (with VLAN and multicast IP address)

```
OS10# show ip igmp snooping groups vlan 3041 232.11.0.0 detail
Interface      vlan3041
Group          232.11.0.0
Source List
 101.41.0.21
   Member Port         Mode                Uptime              Expires
   port-channel51      Include             1d:20:27:36         00:01:09
   ethernet1/1/51:1    Include             1d:20:27:34         00:01:07
   ethernet1/1/52:1    Include             1d:20:27:37         00:01:07
```

Supported Releases 10.4.0E(R1) or later

show ip igmp snooping interface

Displays IGMP snooping interfaces details.

Syntax `show ip igmp snooping interface [vlan vlan-id]`

Parameters *vlan-id*—(Optional) Enter the VLAN ID, from 1 to 4093.

Default Not configured

Command Mode EXEC

Usage Information The multicast flood control feature is not available on the S4248FB-ON, S4248FBL-ON, and S5148-ON devices.

Example

```
OS10# show ip igmp snooping interface
Vlan3031 is up, line protocol is up
IGMP version is 3
IGMP snooping is enabled on interface
IGMP snooping query interval is 60 seconds
IGMP snooping querier timeout is 130 seconds
IGMP snooping last member query response interval is 1000 ms
IGMP Snooping max response time is 10 seconds
IGMP snooping fast-leave is disabled on this interface
IGMP snooping querier is enabled on this interface

Vlan3032 is up, line protocol is up
IGMP version is 3
IGMP snooping is enabled on interface
IGMP snooping query interval is 60 seconds
IGMP snooping querier timeout is 130 seconds
IGMP snooping last member query response interval is 1000 ms
IGMP Snooping max response time is 10 seconds
IGMP snooping fast-leave is disabled on this interface
IGMP snooping querier is enabled on this interface

Vlan3033 is up, line protocol is up
IGMP version is 3
--more--
<<Output Truncated>>
```

```
OS10# show ip igmp snooping interface
Vlan2 is up, line protocol is up
IGMP version is 3
IGMP snooping is enabled on interface
IGMP snooping query interval is 60 seconds
IGMP snooping querier timeout is 130 seconds
IGMP snooping last member query response interval is 1000 ms
IGMP Snooping max response time is 10 seconds
IGMP snooping fast-leave is disabled on this interface
IGMP snooping querier is disabled on this interface
Multicast snooping flood-restrict is enabled on this interface

Vlan3 is up, line protocol is up
IGMP version is 3
IGMP snooping is enabled on interface
IGMP snooping query interval is 60 seconds
IGMP snooping querier timeout is 130 seconds
IGMP snooping last member query response interval is 1000 ms
IGMP Snooping max response time is 10 seconds
IGMP snooping fast-leave is disabled on this interface
IGMP snooping querier is disabled on this interface
Multicast snooping flood-restrict is enabled on this interface
```

Example (with VLAN)

```
OS10# show ip igmp snooping interface vlan 3031
Vlan3031 is up, line protocol is up
IGMP version is 3
IGMP snooping is enabled on interface
IGMP snooping query interval is 60 seconds
IGMP snooping querier timeout is 130 seconds
IGMP snooping last member query response interval is 1000 ms
IGMP Snooping max response time is 10 seconds
IGMP snooping fast-leave is disabled on this interface
IGMP snooping querier is enabled on this interface
```

```
OS10# show ip igmp snooping interface vlan 3031
Vlan3031 is up, line protocol is up
IGMP version is 3
IGMP snooping is enabled on interface
IGMP snooping query interval is 60 seconds
IGMP snooping querier timeout is 130 seconds
IGMP snooping last member query response interval is 1000 ms
IGMP Snooping max response time is 10 seconds
IGMP snooping fast-leave is disabled on this interface
IGMP snooping querier is enabled on this interface
Multicast snooping flood-restrict is enabled on this interface
```

Supported Releases 10.4.0E(R1) or laterUpdated the command to display the multicast flood restrict status on 10.4.3.0 or later

show ip igmp snooping mrouter

Displays the multicast router ports details.

Syntax `show ip igmp snooping mrouter [vlan vlan-id]`

Parameters *vlan-id*—(Optional) Enter the VLAN ID, from 1 to 4093.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip igmp snooping mrouter
Interface   Router Ports
vlan3031    port-channel31
vlan3032    port-channel31
vlan3033    port-channel31
vlan3034    port-channel31
vlan3035    port-channel31
vlan3036    port-channel31
vlan3037    port-channel31
vlan3038    port-channel31
vlan3039    port-channel31
vlan3040    port-channel31
vlan3041    port-channel31
vlan3042    port-channel31
vlan3043    port-channel31
vlan3044    port-channel31
vlan3045    port-channel31
vlan3046    port-channel31
vlan3047    port-channel31
vlan3048    port-channel31
vlan3049    port-channel31
vlan3050    port-channel31
vlan3051    port-channel31
vlan3052    port-channel31
--more--
```

```
<<Output Truncated>>
```

Example (with VLAN)

```
OS10# show ip igmp snooping mrouter vlan 3031
Interface   Router Ports
vlan3031    port-channel31
```

Supported Releases 10.4.0E(R1) or later

Multicast Listener Discovery Protocol

IPv6 networks use Multicast Listener Discovery (MLD) Protocol to manage multicast groups.

OS10 supports MLDv1and MLDv2 to manage the multicast group memberships on IPv6 networks.

MLD snooping

MLD snooping enables switches to use the information in MLD packets and generate a forwarding table that associates ports with multicast groups. When switches receive multicast frames, they forward them to their intended receivers.

OS10 supports MLD snooping on VLAN interfaces. Effective with OS10 release 10.4.3.0, MLD snooping is enabled by default.

Configure MLD snooping

- Enable MLD snooping globally with the `ipv6 mld snooping enable` command in the CONFIGURATION mode. This command enables both MLDv2 and MLDv1 snooping on all VLAN interfaces.
- (Optional) You can disable MLD snooping on specific VLAN interfaces using the `no ipv6 mld snooping` command in the VLAN INTERFACE mode.
- (Optional) Multicast flood control is enabled by default. To disable the multicast flood restrict feature, use the `no multicast snooping flood-restrict` command in CONFIGURATION mode. To reenable the feature globally, use the `ip igmp snooping enable` command in CONFIGURATION mode.
- In a network, the snooping switch is connected to a multicast Router that sends MLD queries. On a Layer 2 network that does not have a multicast router, you can configure the snooping switch to act as querier. Use the `ipv6 mld snooping querier` command in the VLAN INTERFACE mode to send the queries.
- OS10 learns the multicast router interface dynamically based on the interface on which MLD membership query is received. To assign a multicast router interface statically, use the `ipv6 mld snooping mrouter interface` *`interface-type`* command in VLAN INTERFACE mode.
- (Optional) Configure the MLD version using the `ipv6 mld version` *`version-number`* command in the VLAN INTERFACE mode.
- (Optional) The fast leave option allows the MLD snooping switch to remove an interface from the multicast group immediately on receiving the leave message. Enable fast leave with the `ipv6 mld snooping fast-leave` command in VLAN INTERFACE mode.
- (Optional) Configure the time interval for sending MLD general queries with the `ipv6 mld snooping query-interval` *`query-interval-time`* command in VLAN INTERFACE mode.
- (Optional) Configure the maximum time for responding to a query advertised in MLD queries using the `ipv6 mld snooping query-max-resp-time` *`query-response-time`* command in VLAN INTERFACE mode.
- (Optional) Configures the time interval between group-specific MLD query messages with the `ipv6 mld snooping last-member-query-interval` *`query-interval-time`* command in VLAN INTERFACE mode.

MLD snooping configuration

```
OS10(config)# ipv6 mld snooping enable
OS10(config)# interface vlan 11
OS10(conf-if-vl-11)# ipv6 mld snooping mrouter interface ethernet 1/1/32
OS10(conf-if-vl-11)# ipv6 mld snooping querier
OS10(conf-if-vl-11)# ipv6 mld version 1
OS10(conf-if-vl-11)# ipv6 mld snooping fast-leave
OS10(conf-if-vl-11)# ipv6 mld snooping query-interval 60
```

```
OS10(conf-if-vl-11)# ipv6 mld snooping query-max-resp-time 10
OS10(conf-if-vl-11)# ipv6 mld snooping last-member-query-interval 1000
```

View MLD snooping information

```
OS10# show ipv6 mld snooping groups
Total Number of Groups: 280
MLD Connected Group Membership
Group Address                 Interface                     Mode
Expires
ff02::2                       vlan3531                      Exclude
00:01:38
ff0e:225:1::                  vlan3531                      MLDv1-Compat
00:01:52
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::1                 vlan3531                      MLDv1-Compat
00:01:52
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::2                 vlan3531                      MLDv1-Compat
00:01:52
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::3                 vlan3531                      MLDv1-Compat
00:01:52
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::4                 vlan3531                      MLDv1-Compat
00:01:52
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::5                 vlan3531                      MLDv1-Compat
00:01:52
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff02::2                       vlan3532                      Exclude
00:01:47
ff0e:225:2::                  vlan3532                      MLDv1-Compat
00:01:56
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:2::1                 vlan3532                      MLDv1-Compat
00:01:56
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:2::2                 vlan3532                      MLDv1-Compat
00:01:56
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
--more--
<<Output Truncated>>
```

```
OS10# show ipv6 mld snooping interface vlan 3031
Vlan3031 is up, line protocol is up
MLD version is 2
MLD snooping is enabled on interface
MLD snooping query interval is 60 seconds
MLD snooping querier timeout is 130 seconds
MLD snooping last member query response interval is 1000 ms
MLD snooping max response time is 10 seconds
MLD snooping fast-leave is disabled on this interface
MLD snooping querier is disabled on this interface
```

```
OS10# show ipv6 mld snooping interface vlan 2
Vlan2 is up, line protocol is up
MLD version is 2
MLD snooping is enabled on interface
MLD snooping query interval is 60 seconds
MLD snooping querier timeout is 130 seconds
MLD snooping last member query response interval is 1000 ms
MLD snooping max response time is 10 seconds
MLD snooping fast-leave is disabled on this interface
```

```
MLD snooping querier is disabled on this interface
Multicast flood-restrict is enabled on this interface
```

```
OS10# show ipv6 mld snooping mrouter vlan 11
Interface    Router Ports
Vlan 11      ethernet 1/1/32
```

MLD snooping commands

ipv6 mld snooping

Enables MLD snooping on the specified VLAN interface.

Syntax	`ipv6 mld snooping`
Parameters	None
Default	Enabled
Command Mode	VLAN INTERFACE
Usage Information	When you enable MLD snooping globally, the configuration is applied to all the VLAN interfaces. You can disable the MLD snooping on specified VLAN interfaces. The `no` version of this command disables the MLD snooping on the specified VLAN interface.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# no ipv6 mld snooping`
Supported Releases	10.4.1.0 or later

ipv6 mld snooping enable

Enables MLD snooping globally.

Syntax	`ipv6 mld snooping enable`
Parameters	None
Default	Enabled
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command disables the MLD snooping.
Example	`OS10(config)# ipv6 mld snooping enable`
Supported Releases	10.4.1.0 or later

ipv6 mld snooping fast-leave

Enables fast leave in MLD snooping for specified VLAN.

Syntax	`ipv6 mld snooping fast-leave`
Parameters	None
Default	Disabled

Command Mode	VLAN INTERFACE
Usage Information	The fast leave option allows the MLD snooping switch to remove an interface from the multicast group immediately on receiving the `leave` message. The `no` version of this command disables the fast leave functionality.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# ipv6 mld snooping fast-leave`
Supported Releases	10.4.1.0 or later

ipv6 mld snooping last-member-query-interval

Configures the time interval between group-specific MLD query messages.

Syntax	`ipv6 mld snooping last-member-query-interval` *`query-interval-time`*
Parameters	*`query-interval-time`*—Enter the query time interval in milliseconds, ranging from 100 to 65535.
Default	1000 milliseconds
Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command resets the last member query interval time to the default value.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# ipv6 mld snooping last-member-query-interval 2500`
Supported Releases	10.4.1.0 or later

ipv6 mld snooping mrouter

Configures the specified VLAN member port as a multicast router interface.

Syntax	`ipv6 mld snooping mrouter interface` *`interface-type`*
Parameters	*`interface-type`*—Enter the interface type details. The interface should be a member of the VLAN.
Default	Not configured
Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command removes the multicast router configuration from the VLAN member port.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# ipv6 mld snooping mrouter interface ethernet 1/1/1`
Supported Releases	10.4.1.0 or later

ipv6 mld snooping querier

Enables MLD querier on the specified VLAN interface.

Syntax	`ipv6 mld snooping querier`
Parameters	None
Default	Not configured

Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command disables the MLD querier on the VLAN interface.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# ipv6 mld snooping querier`
Supported Releases	10.4.1.0 or later

ipv6 mld snooping query-interval

Configures the time interval for sending MLD general queries.

Syntax	`ipv6 mld snooping query-interval` *`query-interval-time`*
Parameters	*`query-interval-time`*—Enter the interval time in seconds, ranging from 2 to 18000.
Default	60 seconds
Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command resets the query interval to the default value.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# ipv6 mld snooping query-interval 120`
Supported Releases	10.4.1.0 or later

ipv6 mld query-max-resp-time

Configures the maximum time for responding to a query advertised in MLD queries.

Syntax	`ipv6 mld snooping query-max-resp-time` *`query-response-time`*
Parameters	*`query-response-time`*—Enter the query response time in seconds, ranging from 1 to 25.
Default	10 seconds
Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command resets the query response time to default value.
Example	`OS10(config)# interface vlan 100` `OS10(conf-if-vl-100)# ipv6 mld snooping query-max-resp-time 15`
Supported Releases	10.4.1.0 or later

ipv6 mld version

Configures the MLD version.

Syntax	`ipv6 mld version` *`version-number`*
Parameters	*`version-number`*—Enter the version number as 1 or 2.
Default	2
Command Mode	VLAN INTERFACE
Usage Information	The `no` version of this command resets the version number to the default value.

Example

```
OS10(config)# interface vlan 100
OS10(conf-if-vl-100)# ipv6 mld version 1
```

Supported Releases 10.4.1.0 or later

show ipv6 mld snooping groups

Displays MLD snooping group membership details.

Syntax `show ipv6 mld snooping groups [vlan vlan-id] [ipv6-address]`

Parameters

- *vlan-id*—(Optional) Enter the VLAN ID, from 1 to 4093.
- *ipv6-address*—(Optional) Enter the IPv6 address of the multicast group.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show ipv6 mld snooping groups
Total Number of Groups: 280
MLD Connected Group Membership
Group Address                  Interface                      Mode
Expires
ff02::2                        vlan3531                       Exclude
00:01:38
ff0e:225:1::                   vlan3531                       MLDv1-Compat
00:01:52
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::1                  vlan3531                       MLDv1-Compat
00:01:52
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::2                  vlan3531                       MLDv1-Compat
00:01:52
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::3                  vlan3531                       MLDv1-Compat
00:01:52
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::4                  vlan3531                       MLDv1-Compat
00:01:52
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::5                  vlan3531                       MLDv1-Compat
00:01:52
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff02::2                        vlan3532                       Exclude
00:01:47
ff0e:225:2::                   vlan3532                       MLDv1-Compat
00:01:56
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:2::1                  vlan3532                       MLDv1-Compat
00:01:56
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:2::2                  vlan3532                       MLDv1-Compat
00:01:56
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
--more--
```

Example (with VLAN)

```
OS10# show ipv6 mld snooping groups vlan 3531
Total Number of Groups: 7
MLD Connected Group Membership
Group Address                  Interface           Mode          Expires
ff02::2                        vlan3531            Exclude       00:02:08
```

```
ff0e:225:1::                    vlan3531            MLDv1-Compat  00:02:12
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::1                   vlan3531            MLDv1-Compat  00:02:12
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::2                   vlan3531            MLDv1-Compat  00:02:12
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::3                   vlan3531            MLDv1-Compat  00:02:12
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::4                   vlan3531            MLDv1-Compat  00:02:12
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
ff0e:225:1::5                   vlan3531            MLDv1-Compat  00:02:12
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
```

Example (with VLAN and multicast IP address)

```
OS10# show ipv6 mld snooping groups vlan 3531 ff0e:225:1::
MLD Connected Group Membership
Group Address      Interface           Mode          Expires
ff0e:225:1::      vlan3531            MLDv1-Compat  00:01:30
  Member-ports :port-channel41,ethernet1/1/51,ethernet1/1/52
```

Supported Releases 10.4.0E(R1) or later

show ipv6 mld snooping groups detail

Displays the MLD source information along with detailed member port information.

Syntax `show ipv6 mld snooping groups [vlan vlan-id] [group ipv6-address] detail`

Parameters

- *vlan-id*—(Optional) Enter the VLAN ID, ranging from 1 to 4093.
- *ipv6-address*—(Optional) Enter the IPv6 address of the multicast group.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show ipv6 mld snooping groups detail
Interface      vlan3041
Group          ff02::2
Source List
 --
   Member Port         Mode                Uptime              Expires
   port-channel31      Exclude             2d:11:57:08         00:01:44

Interface      vlan3041
Group          ff3e:232:b::
Source List
 2001:101:29::1b
   Member Port         Mode                Uptime              Expires
   port-channel31      Include             2d:11:50:17         00:01:42
   ethernet1/1/51:1    Include             2d:11:50:36         00:01:38
   ethernet1/1/52:1    Include             2d:11:50:36         00:01:25

Interface      vlan3041
Group          ff3e:232:b::1
Source List
 2001:101:29::1b
   Member Port         Mode                Uptime              Expires
   port-channel31      Include             2d:11:50:17         00:01:29
   ethernet1/1/51:1    Include             2d:11:50:36         00:01:25
   ethernet1/1/52:1    Include             2d:11:50:36         00:01:38
--more--
```

Example (with VLAN)

```
OS10# show ipv6 mld snooping groups vlan 3041 detail
Interface      vlan3041
Group          ff02::2
Source List
 --
   Member Port         Mode               Uptime              Expires
   port-channel31      Exclude            2d:11:57:08         00:01:44

Interface      vlan3041
Group          ff3e:232:b::
Source List
 2001:101:29::1b
   Member Port         Mode               Uptime              Expires
   port-channel31      Include            2d:11:50:17         00:01:42
   ethernet1/1/51:1    Include            2d:11:50:36         00:01:38
   ethernet1/1/52:1    Include            2d:11:50:36         00:01:25

Interface      vlan3041
Group          ff3e:232:b::1
Source List
 2001:101:29::1b
   Member Port         Mode               Uptime              Expires
   port-channel31      Include            2d:11:50:17         00:01:29
   ethernet1/1/51:1    Include            2d:11:50:36         00:01:25
   ethernet1/1/52:1    Include            2d:11:50:36         00:01:38
--more--
```

Example (with VLAN and multicast IP address)

```
OS10# show ipv6 mld snooping groups vlan 3041 ff3e:232:b:: detail
Interface      vlan3041
Group          ff3e:232:b::
Source List
 2001:101:29::1b
   Member Port         Mode               Uptime              Expires
   port-channel31      Include            2d:11:50:53         00:02:01
   ethernet1/1/51:1    Include            2d:11:51:11         00:02:01
   ethernet1/1/52:1    Include            2d:11:51:12         00:01:52
```

Supported Releases 10.4.1.0 or later

show ipv6 mld snooping interface

Displays the details of MLD snooping interfaces.

Syntax `show ipv6 mld snooping interface [vlan vlan-id]`

Parameters *vlan-id*—(Optional) Enter the VLAN ID, ranging from 1 to 4093.

Default Not configured

Command Mode EXEC

Usage Information The multicast flood control feature is not available on the S4248FB-ON, S4248FBL-ON, and S5148-ON devices.

Example

```
OS10# show ipv6 mld snooping interface vlan 3031
Vlan3031 is up, line protocol is up
MLD version is 2
MLD snooping is enabled on interface
MLD snooping query interval is 60 seconds
MLD snooping querier timeout is 130 seconds
MLD snooping last member query response interval is 1000 ms
MLD snooping max response time is 10 seconds
```

```
MLD snooping fast-leave is disabled on this interface
MLD snooping querier is disabled on this interface
```

```
OS10# show ipv6 mld snooping interface vlan 2
Vlan2 is up, line protocol is up
MLD version is 2
MLD snooping is enabled on interface
MLD snooping query interval is 60 seconds
MLD snooping querier timeout is 130 seconds
MLD snooping last member query response interval is 1000 ms
MLD snooping max response time is 10 seconds
MLD snooping fast-leave is disabled on this interface
MLD snooping querier is disabled on this interface
Multicast flood-restrict is enabled on this interface
```

Supported Releases 10.4.1.0 or laterUpdated the command to display the multicast flood restrict status on 10.4.3.0 or later

show ipv6 mld snooping mrouter

Displays the details of multicast router ports.

Syntax `show ipv6 mld snooping mrouter [vlan vlan-id]`

Parameters *vlan-id*—(Optional) Enter the VLAN ID, ranging from 1 to 4093.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show ipv6 mld snooping mrouter vlan 11
Interface    Router Ports
Vlan 11      ethernet 1/1/32
```

Supported Releases 10.4.1.0 or later

Protocol Independent Multicast

Protocol independent multicast (PIM) is a group of multicast routing protocols that provides one-to-many and many-to-many transmission of information. PIM uses routing information from other routing protocols and does not depend on any specific unicast routing protocol. PIM uses any unicast routing protocol that is deployed in the network. OS10 supports the following PIM modes:

- PIM sparse mode (PIM-SM)
- PIM source specific multicast (PIM-SSM)

PIM terminology

Table 16. PIM terminology

Terminology	Definition
Rendezvous point (RP)	The RP is a single root node that the shared tree uses, called the rendezvous point.
(*, G)	(*, G) refers to an entry in the PIM table for a group.
(S, G)	(S, G) refers to an entry in the PIM table for a source and group on the RP tree (RPT).

Terminology	Definition
(S, G, RPT)	(S, G, RPT) refers to an entry in the RP tree.
First hop router (FHR)	The FHR is the router that is directly connected to the multicast source.
Last hop router (LHR)	The LHR is the last router in the multicast path and is directly connected to the multicast receiver.
Intermediate router	A PIM router that is not an FHR, RP, or LHR.
Shared tree (RPT)	The RPT is an unidirectional multicast tree whose root node is the RP.
Shortest path tree (SPT)	The root node of the SPT is the multicast source. The multicast traffic routes to the receiver on the shortest path. This setup reduces network latency and traffic congestion at the RP.
Outgoing interface (OIF)	The OIF is the interface through which a multicast packet is sent out towards the receiver.
Incoming interface (IIF)	The IIF is the interface through which a multicast packet is received towards the source or the RP.
Reverse path forwarding (RPF)	The RPF is the path the router uses to reach the RP or the multicast source.

Standards compliance

OS10 complies to the following standards:

- RFC 4601 for PIM-SM
- RFC 3569 for PIM-SSM

PIM-SM

PIM sparse mode (PIM-SM) is a multicast routing protocol for networks with receivers that are sparsely distributed. Receivers have to explicitly send a *join* message to join particular groups or sources. PIM join and prune messages are used to join and leave multicast distribution trees.

PIM-SM uses shared trees with the root node being the rendezvous point (RP). All multicast sources use the RP to route the traffic to the receiver. The last hop router (LHR) sends an (*,G) join message towards the RP. The designated router connected to the first hop router (FHR) encapsulates multicast data that comes from the multicast source in PIM control messages and sends it via unicast to the RP as PIM register messages. The RP sends an (S, G) join towards the source. When the RP receives native data traffic from the source, it sends a register stop message to the FHR.

OS10 supports static configuration of an RP address for a multicast group.

To keep the PIM-SM state alive, all PIM join messages are periodically re-transmitted.

You must enable PIM-SM on each of the participating interfaces. Be sure to have multicast routing enabled on the system. To do this, use the `ip multicast-routing` command from CONFIGURATION mode.

```
OS10# configure terminal
OS10(config)# interface vlan 100
OS10(conf-if-vl-100)# ip pim sparse-mode
```

PIM-SSM

PIM-SSM uses source-based trees. A separate multicast distribution tree is built for each multicast source that sends data to a multicast group. Each multicast distribution tree has as its root node a router adjacent to the source. Sources send data directly to the root of the tree. PIM-SSM allows receivers to specify the source from which to receive data as well as the multicast group they want to join. The receiver identifies a multicast data stream using the source and group address pair (S, G) instead of the group address alone (*, G).

NOTE:

- PIM-SSM requires receivers to support IGMP version 3.
- The default PIM-SSM range is 232.0.0.0/8. The default range is always supported and the range can never be smaller than the default.

Advantages of PIM-SSM

Advantages of PIM-SSM include the following:

- PIM-SSM forwards multicast traffic from a single source to a subnet. Other versions of PIM requires the receiver to subscribe to a group. The receiver receives traffic not just from the source that it is interested in, but from all the sources that send to that group. PIM-SSM requires the receiver to specify the sources in which they are interested in to avoid receiving unnecessary traffic.
- PIM-SSM is more efficient than PIM-SM because it immediately creates shortest path trees (SPT) to the source rather than using shared trees. PIM-SM requires a shared tree rooted at the RP because IGMPv2 receivers do not express the source information in their membership reports. Multicast traffic passes from the source to the receiver through the RP, until the last hop router (LHR) learns the source address, at which point it switches to the SPT.

  PIM-SSM uses IGMPv3. Because receivers subscribe to a source and group, the RP and shared tree are unnecessary; only SPTs are used. On OS10 systems, it is possible to use PIM-SM with IGMPv3 to achieve the same result, but PIM-SSM eliminates the unnecessary protocol overhead.

Configure PIM-SSM

To configure a group range for PIM-SSM:

NOTE: The IP range, 232.0.0.0/8 is reserved for SSM. You do not have to explicitly configure this range.

1 Create an ACL rule to specify the range of addresses that should use SSM.

```
OS10# configure terminal
OS10(config)# ip access-list ssm-1
OS10(config-ipv4-acl)# permit ip any 236.0.0.0/8
OS10(config-ipv4-acl)# exit
```

2 Enable PIM-SSM for the range of addresses using the `ip pim ssm-range` command.

```
OS10(config)# ip pim ssm-range ssm-1
```

You can use the `show ip pim ssm-range` command to view the groups added in PIM-SSM configuration.

```
OS10# show ip pim ssm-range
Group Address    / MaskLen
236.0.0.0        / 8
```

Expiry timers for S, G entries

You can configure expiry timers for S, G entries globally. The S, G entries expire in 210 seconds by default.

To configure the S, G expiry timer:

```
OS10# configure terminal
OS10(config)# ip pim sparse-mode sg-expiry-timer 500
```

Static rendezvous point

The rendezvous point (RP) is an interface on a router that acts as the root to a group-specific tree; every group must have an RP. You must configure the RP on all nodes in your network.

To configure a static RP:

```
OS10# configure terminal
OS10(config)# ip pim rp-address 171.1.1.1 group-address 225.1.1.3/32
```

Designated router

Multiple PIM-SM routers can connect to a single local area network (LAN) segment. One of these routers is elected as the designated router (DR).

The DR is elected using hello messages. Each PIM router learns about its neighbors by periodically sending a hello message out of each PIM-enabled interface. Hello messages contain the IP address of the interface from where it is sent and a DR priority value. The router with the greatest priority value becomes the DR. If the priority value is the same for two routers, the router with the highest IP address is the DR. By default, the DR priority value is 1, so the IP address determines the DR.

To configure DR priority, use the following command:

```
OS10# configure terminal
OS10(config)# interface vlan 100
OS10(conf-if-vl-100)# ip pim dr-priority 200
```

PIM commands

clear ip pim tib

Clears PIM tree information from the PIM database.

Syntax	`clear ip pim [vrf vrf-name] tib`
Parameters	`vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.
Default	None
Command Mode	EXEC PRIVILEGE
Usage Information	When you run this command on a node, all multicast routes from the PIM tree information base (TIB), the entire multicast route table, and all the entries in the data plane are deleted.
Example	`OS10# clear ip pim vrf vrf1 tib`
Supported Releases	10.4.3.0 or later

ip multicast-routing

Enables IP multicast forwarding.

Syntax	`ip multicast-routing [vrf vrf-name]`
Parameters	`vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF to enable IP multicast forwarding on the specified VRF.
Default	None
Command Mode	CONFIGURATION
Usage Information	After you enable IP multicast, enable IGMP and PIM on an interface. To do this, use the `ip pim sparse-mode` command in INTERFACE mode. The `no` form of the command disables IP multicast forwarding.
Example	`OS10# configure terminal` `OS10(config)# ip multicast-routing`
Supported Releases	10.4.3.0 or later

ip pim dr-priority

Changes the designated router (DR) priority for the interface.

Syntax	`ip pim dr-priority priority-value`
Parameters	`priority-value`—Enter a number from 0 to 4294967295.
Default	1
Command Mode	INTERFACE CONFIGURATION
Usage Information	The router with the largest value assigned to an interface becomes the DR. If two interfaces have the same DR priority value, the interface with the largest IP address becomes the DR. Use the `no` form of the command to remove the DR priority value assigned to an interface.

Example	`OS10# configure terminal`
	`OS10(config)# interface vlan 1`
	`OS10(conf-if-vl-1)# ip pim dr-priority 200`

Supported Releases	10.4.3.0 or later

ip pim query-interval

Changes the frequency of PIM router query messages.

Syntax	`ip pim query-interval seconds`
Parameters	`seconds`—Enter the amount of time, in seconds, the router waits before sending a PIM hello packet out of each PIM-enabled interface, from 2 to 18000.
Default	30 seconds
Command Mode	INTERFACE CONFIGURATION

Usage Information Use the `no` form of the command to return the frequency of PIM router query messages to the default value.

Example

```
OS10# configure terminal
OS10(config)# interface vlan 1
OS10(conf-if-vl-1)# ip pim query-interval 20
```

Supported Releases 10.4.3.0 or later

ip pim rp-address

Configures a static PIM RP address for a group.

Syntax `ip pim [vrf vrf-name] rp-address address {group-address group-address mask}`

Parameters

- `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.
- `rp-address address`—Enter the keyword `address`, then the RP address in dotted-decimal format (A.B.C.D).
- `group-address group-address mask`—Enter the keyword `group-address`, then the group-address mask in dotted-decimal format (/xx) to assign the group address to the RP.

Default None

Command Mode CONFIGURATION

Usage Information First hop routers use this address to send register packets on behalf of the source multicast hosts. The RP addresses are stored in the order in which they are entered. The RP is chosen based on a longer prefix match for a group.

Example

```
OS10# configure terminal
OS10(config)# ip pim rp-address 171.1.1.1 group-address 225.1.1.3/32
```

Supported Releases 10.4.3.0 or later

ip pim sparse-mode

Enables PIM sparse mode and IGMP on the interface.

Syntax `ip pim sparse-mode`

Parameters None

Default Disabled

Command Mode INTERFACE CONFIGURATION

Usage Information Before you enable PIM sparse mode, ensure that:

- Multicast is enabled globally using the `ip multicast-routing` command.
- The interface is enabled. Use the `no shutdown` command to enable the interface.
- The interface is in Layer 3 mode. PIM-SM is enabled only on a Layer 3 interface. Before configuring PIM on the interface, use the `no switchport` command to change the interface from Layer 2 to Layer 3 mode.

Use the `no` form of the command to disable PIM sparse mode.

Example

```
OS10# configure terminal
OS10(config)# interface vlan 2
```

```
OS10(conf-if-vl-2)# ip address 1.1.1.2/24
OS10(conf-if-vl-2)# ip pim sparse-mode
```

Supported Releases 10.4.3.0 or later

ip pim sparse-mode sg-expiry-timer

Enables expiry timers globally for all sources.

Syntax `ip pim [vrf vrf-name] sparse-mode sg-expiry-timer seconds`

Parameters

- `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.
- *`seconds`*—Enter the number of seconds the S, G entries are retained. The range is from 211 to 65535 seconds.

Default 210 seconds

Command Mode CONFIGURATION

Usage Information This command configures the expiry timers for all S, G entries.

Example

```
OS10# configure terminal
OS10(config)# ip pim sparse-mode sg-expiry-timer 500
```

Supported Releases 10.4.3.0 or later

ip pim ssm-range

Specifies the SSM group range using an access list.

Syntax `ip pim [vrf vrf-name] ssm-range {access-list-name}`

Parameters

- `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.
- *`access-list-name`*—Enter the name of the access list.

Default 232.0.0.0/8

Command Mode CONFIGURATION

Usage Information When ACL rules change, the ACL and PIM modules apply the new rules automatically. When you remove the SSM ACL, PIM-SSM is supported only for the default SSM range.

Example

```
OS10# configure terminal
OS10(config)# ip pim ssm-range ssm
```

Supported Releases 10.4.3.0 or later

show ip pim interface

Displays information about IP PIM-enabled interfaces.

Syntax `show ip pim [vrf vrf-name] interface`

Parameters `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.

Default None

Command Mode EXEC

Usage Information The `show ip pim interface` command displays the following:

- `Address`—IP addresses of the IP PIM-enabled interfaces
- `Interface`—Interface type with slot/port information or VLAN/Port Channel ID
- `Version/Mode`—PIM version number and mode; `v2` for PIM version 2 and `S` for PIM sparse mode
- `Nbr Count`—Active neighbor count on the PIM-enabled interface
- `Query interval`—Query interval for router query messages on that interface
- `DR priority`—Designated router priority value configured on that interface
- `DR`—IP address of the DR for that interface

Example

```
OS10# show ip pim interface
Address         Interface            Ver/Mode       Nbr Count    Query Intvl      DR Prio      DR
-------------------------------------------------------------------------------------------------------
2.2.2.2         vlan103              v2/S           1            30               1            2.2.2.2
3.3.3.2         vlan105              v2/S           1            30               1            3.3.3.2
122.1.1.2       vlan121              v2/S           0            30               1            122.1.1.2
```

Supported Releases 10.4.3.0 or later

show ip pim mcache

Displays routes synchronized from VLT peer and local route information.

Syntax `show ip pim [vrf vrf-name] mcache [group-address [source-address]]`

Parameters

- `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.
- *`group-address`*—Enter the multicast group address in dotted-decimal format (A.B.C.D).
- *`source-address`*—Enter the multicast source address in dotted-decimal format (A.B.C.D).

Default None

Command Mode EXEC

Usage Information This command provides details about the incoming and outgoing interfaces for multicast routes.

Example

```
OS10# show ip pim mcache
PIM Multicast Routing Cache Table

(*, 225.1.1.1)
  Incoming interface : vlan105
  Outgoing interface list :
   vlan121

(101.1.1.10,225.1.1.1)
  Incoming interface : vlan103
  Outgoing interface list :
   vlan121
```

Supported Releases 10.4.3.0 or later

show ip pim neighbor

Displays PIM neighbors.

Syntax `show ip pim [vrf vrf-name] neighbor`

Parameters `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.

Default None

Command Mode EXEC

Usage Information This command displays the following:

- `Neighbor address`—IP addresses of the PIM neighbor
- `Interface`—Interface type with slot/port information or VLAN/Port Channel ID of the PIM neighbor
- `Uptime/expires`—Amount of time that the PIM neighbor has been up
- `Version`—PIM version number; `v2` for PIM version 2
- `DR priority/Mode`—Designated router priority value and mode. The default designated router priority is 1 and S for sparse mode

Example

```
OS10# show ip pim neighbor
Neighbor Address Interface  Uptime/Expires         Ver         DR Prio/Mode
--------------------------------------------------------------------------------
2.1.1.1          vlan103    13:05:58/00:01:19        v2           1   / S
3.1.1.1          vlan105    13:05:58/00:01:17        v2           1   / S
```

Supported Releases 10.4.3.0 or later

show ip pim rp

Displays brief information about all multicast group to RP mappings.

Syntax `show ip pim [vrf vrf-name] rp [mapping | group-address]`

Parameters

- `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.
- `mapping`—Enter the keyword `mapping` to display the multicast groups to RP mapping and information about how RP is learned.
- *`group-address`*—Enter the multicast group address mask in dotted-decimal format to view the RP for a specific group (A.B.C.D).

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip pim rp
Group           RP
---------------------------------
225.1.1.1       171.1.1.1
225.1.1.2       171.1.1.1
225.1.1.3       171.1.1.1
225.1.1.4       171.1.1.1
225.1.1.5       171.1.1.1
225.1.1.6       171.1.1.1
```

```
225.1.1.7       171.1.1.1
225.1.1.8       171.1.1.1
225.1.1.9       171.1.1.1
225.1.1.10      171.1.1.1
225.1.1.11      171.1.1.1
225.1.1.12      171.1.1.1
225.1.1.13      171.1.1.1
```

Supported Releases 10.4.3.0 or later

show ip pim ssm-range

Displays the non-default groups added using the SSM range feature.

Syntax `show ip pim [vrf vrf-name] ssm-range`

Parameters `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip pim ssm-range
Group Address   / MaskLen
224.1.1.1       / 32
```

Supported Releases 10.4.3.0 or later

show ip pim summary

Displays information about PIM-SM operation.

Syntax `show ip pim [vrf vrf-name] summary`

Parameters `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip pim summary

Entries in PIM-TIB/MFC: 300/200

Active Modes:
    PIM-SM

Interface summary:
    2 active PIM interfaces
    0 passive PIM interfaces
    1 active PIM neighbor

TIB Summary:
    100/100 (*,G) entries in PIM-TIB/MFC
    100/100 (S,G) entries in PIM-TIB/MFC
    100/0 (S,G,Rpt) entries in PIM-TIB/MFC

    1 RP
```

```
    1 sources

Message Summary:
    0/0 Joins/Prunes sent/received
    294 Null Register messages received
    404/404 Register-stop messages sent/received

Data path event summary:
    0/0 pim-assert messages sent/received
    404/110 register messages sent/received
```

Supported Releases 10.4.3.0 or later

show ip pim tib

Displays the PIM tree information base (TIB).

Syntax `show ip pim [vrf vrf-name] tib [group-address [source-address]]`

Parameters

- `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.
- `group-address`—Enter the group address in dotted-decimal format (A.B.C.D).
- `source-address`—Enter the source address in dotted-decimal format (A.B.C.D).

Default None

Command Mode EXEC

Usage Information This command displays the following:

- `S, G`—Displays the entry in the multicast PIM database
- `uptime`—Displays the amount of time the entry has been in the PIM route table
- `expires`—Displays the amount of time until the entry expires and is removed from the database
- `RP`—Displays the IP address of the RP or source for the entry
- `Incoming interface`—Displays the reverse path forwarding (RPF) interface towards the RP/source
- `RPF neighbor`—Displays the next hop IP address from this interface towards the RP/source
- `Outgoing interface list`—Lists the interfaces that meet one of the following criteria:
  - a directly connected member of the group
  - a statically connected member of the group
  - received an (*, G) or (S, G) join message

Example

```
OS10# show ip pim tib

PIM Multicast Routing Table
Flags: S - Sparse, C - Connected, L - Local, P - Pruned,
       R - RP-bit set, F - Register Flag, T - SPT-bit set, J - Join SPT,
       K - Ack-Pending state
Timers: Uptime/Expires
Interface state: Interface, next-Hop, State/Mode

(*, 225.1.1.1), uptime 13:08:24, expires 00:00:12, RP 171.1.1.1, flags: SCJ
  Incoming interface: vlan105, RPF neighbor 3.1.1.1
  Outgoing interface list:
    vlan121  Forward/Sparse  13:07:53/Never


(101.1.1.10, 225.1.1.1), uptime 13:07:51, expires 00:06:09, flags: T
  Incoming interface: vlan103, RPF neighbor 2.1.1.1
```

```
  Outgoing interface list:
    vlan121  Forward/Sparse  13:07:50/Never
```

Supported Releases 10.4.3.0 or later

show ip rpf

Displays reverse path forwarding (RPF) information.

Syntax `show ip rpf [vrf vrf-name] [source-address] [summary]`

Parameters

- `vrf vrf-name`—Enter the keyword `vrf`, then the name of the VRF.
- `source-address`—Enter the source address in dotted-decimal format (A.B.C.D).
- `summary`—RPF summary.

Default None

Command Mode EXEC

Usage Information Use static mroutes to control the reachability of the multicast sources. If a PIM-registered multicast source is reachable using static mroute as well as a unicast route, PIM examines the distance of each route. The route with shorter distance is the one that PIM selects for reachability.

Example

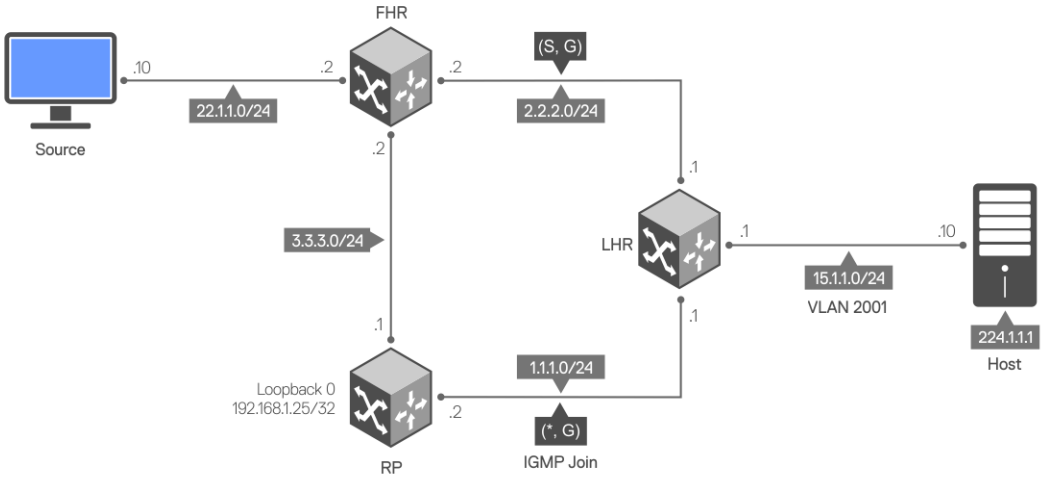
```
OS10# show ip rpf
RPF information for 101.1.1.10
   RPF interface: vlan103
   RPF neighbor: 2.1.1.1
   RPF route/mask: 101.1.1.0/255.255.255.0
   RPF type: Unicast
RPF information for 171.1.1.1
   RPF interface: vlan105
   RPF neighbor: 3.1.1.1
   RPF route/mask: 171.1.1.0/255.255.255.0
   RPF type: Unicast
```

Supported Releases 10.4.3.0 or later

PIM-SM sample configuration

This section describes how to enable PIM-SM in the FHR, RP, and LHR nodes using the topology show in the following illustration.
To enable PIM-SM, perform the following configurations on each of the nodes (FHR, RP, and LHR):

- Enable multicast routing on all the nodes using the `ip multicast-routing` command
- Enable PIM-SM on all the required Layer 3 interfaces of the nodes using the `ip pim sparse-mode` command
- Configure an RP address on every multicast enable node using the `ip pim rp-address` command
- Configure an IP address for each interface of the nodes in the PIM-SM topology
- Enable a routing protocol (OSPF) for route updates



Sample configuration in FHR node:

```
FHR# configure terminal
FHR(config)#
FHR(config)# ip multicast-routing
FHR(config)# interface ethernet 1/1/31
FHR(conf-if-eth1/1/31)# no switchport
FHR(conf-if-eth1/1/31)# ip address 3.3.3.2/24
FHR(conf-if-eth1/1/31)# ip pim sparse-mode
FHR(conf-if-eth1/1/31)# ip ospf 1 area 0
FHR(conf-if-eth1/1/31)# exit
FHR(config)#
FHR(config)# interface ethernet 1/1/17
FHR(conf-if-eth1/1/17)#
FHR(conf-if-eth1/1/17)# no switchport
FHR(conf-if-eth1/1/17)# ip address 2.2.2.2/24
FHR(conf-if-eth1/1/17)# ip pim sparse-mode
FHR(conf-if-eth1/1/17)# ip ospf 1 area 0
FHR(conf-if-eth1/1/17)# exit
FHR(config)# router ospf 1
FHR(config-router-ospf-1)# exit
FHR(config)# ip pim rp-address 192.168.1.25 group-address 224.0.0.0/4
FHR(config)#
FHR# configure terminal
FHR(config)# interface ethernet 1/1/48
FHR(conf-if-eth1/1/48)# no switchport
FHR(conf-if-eth1/1/48)# ip address 22.1.1.2/24
FHR(conf-if-eth1/1/48)# ip pim sparse-mode
FHR(conf-if-eth1/1/48)# ip ospf 1 area 0
FHR(conf-if-eth1/1/48)#
```

The `show ip pim interface` command displays the PIM-enabled interfaces in FHR.

```
FHR# show ip pim interface
Address       Interface           Ver/Mode      Nbr Count    Query Intvl      DR Prio      DR
-------------------------------------------------------------------------------------------
2.2.2.2       ethernet1/1/17      v2/S          1            30               1            2.2.2.2
3.3.3.2       ethernet1/1/31      v2/S          1            30               1            3.3.3.2
22.1.1.2      ethernet1/1/48      v2/S          0            30               1            22.1.1.2
FHR#
```

The `show ip pim neighbor` command displays the PIM neighbor of FHR and the interface to reach the neighbor.

```
FHR# show ip pim neighbor
Neighbor Address          Interface             Uptime/Expires         Ver         DR Prio/Mode
-----------------------------------------------------------------------------------------------
2.2.2.1                   ethernet1/1/17      00:04:31/00:01:43         v2            1    / S
3.3.3.1                   ethernet1/1/31      00:05:45/00:01:31         v2            1    / S
FHR#
```

The `show ip pim rp mapping` command displays the multicast groups to RP mapping and information about how RP is learned.

```
FHR# show ip pim rp mapping
Group(s) : 224.0.0.0/4,  Static
RP : 192.168.1.25, v2
```

Sample configuration in RP node:

```
RP# configure terminal
RP(config)# ip multicast-routing
RP(config)# interface ethernet 1/1/31
RP(conf-if-eth1/1/31)# no switchport
RP(conf-if-eth1/1/31)# ip address 3.3.3.1/24
RP(conf-if-eth1/1/31)# ip pim sparse-mode
RP(conf-if-eth1/1/31)# ip ospf 1 area 0
RP(conf-if-eth1/1/31)# exit
RP(config)#
RP(config)# interface ethernet 1/1/43
RP(conf-if-eth1/1/43)# no switchport
RP(conf-if-eth1/1/43)# ip address 1.1.1.2/24
RP(conf-if-eth1/1/43)# ip pim sparse-mode
RP(conf-if-eth1/1/43)# ip ospf 1 area 0
RP(conf-if-eth1/1/43)# exit
RP(config)#
RP(config)# interface loopback 0
RP(conf-if-lo-0)# ip address 192.168.1.25/32
RP(conf-if-lo-0)# ip ospf 1 area 0
RP(conf-if-lo-0)# exit
RP(config)# ip pim rp-address 192.168.1.25 group-address 224.0.0.0/4
RP(config)# end
RP#
RP# configure terminal
RP(config)# router ospf 1
RP(config-router-ospf-1)# end
```

The `show ip pim interface` command displays the PIM-enabled interfaces in RP.

```
RP# show ip pim interface
Address      Interface            Ver/Mode       Nbr Count    Query Intvl      DR Prio      DR
-------------------------------------------------------------------------------------------
3.3.3.1      ethernet1/1/31       v2/S           1            30               1            3.3.3.2
1.1.1.2      ethernet1/1/43       v2/S           1            30               1            1.1.1.2
RP#
```

The `show ip pim neighbor` command displays the PIM neighbor of RP and the interface to reach the neighbor.

```
RP# show ip pim neighbor
Neighbor Address          Interface             Uptime/Expires          Ver         DR Prio/Mode
-------------------------------------------------------------------------------------------
3.3.3.2                   ethernet1/1/31      00:02:57/00:01:17         v2            1   / DR S
1.1.1.1                   ethernet1/1/43      00:06:35/00:01:39         v2            1   / S
RP#
```

Sample configuration in LHR node:

```
LHR# configure terminal
LHR(config)# ip multicast-routing
LHR(config)# interface ethernet 1/1/17
LHR(conf-if-eth1/1/17)#
LHR(conf-if-eth1/1/17)# no switchport
LHR(conf-if-eth1/1/17)# ip address 1.1.1.1/24
LHR(conf-if-eth1/1/17)# ip pim sparse-mode
LHR(conf-if-eth1/1/17)# ip ospf 1 area 0
LHR(conf-if-eth1/1/17)# exit
LHR(config)#
LHR(config)# interface ethernet 1/1/29
LHR(conf-if-eth1/1/29)# no switchport
LHR(conf-if-eth1/1/29)# ip address 2.2.2.1/24
LHR(conf-if-eth1/1/29)# ip pim sparse-mode
LHR(conf-if-eth1/1/29)# ip ospf 1 area 0
```

```
LHR(conf-if-eth1/1/29)# exit
LHR(config)#
LHR(config)# ip pim rp-address 192.168.1.25 group-address 224.0.0.0/4
LHR(config)# end
LHR(config)# interface vlan 2001
LHR(conf-if-vl-2001)# no shutdown
LHR(conf-if-vl-2001)# ip address 15.1.1.1/24
LHR(conf-if-vl-2001)# ip pim sparse-mode
LHR(conf-if-vl-2001)# ip ospf 1 area 0
LHR(conf-if-vl-2001)# exit
LHR(config)#
LHR(config)# interface ethernet 1/1/38
LHR(conf-if-eth1/1/38)# switchport mode trunk
LHR(conf-if-eth1/1/38)# no switchport access vlan
LHR(conf-if-eth1/1/38)# switchport trunk allowed vlan 2001
LHR(conf-if-eth1/1/38)# exit
LHR# configure terminal
LHR(config)# router ospf 1
LHR(config-router-ospf-1)# end
```

The `show ip pim interface` command displays the PIM-enabled interfaces in LHR.

```
LHR# show ip pim interface
Address         Interface            Ver/Mode       Nbr Count    Query Intvl      DR Prio
DR
-------------------------------------------------------------------------------------------------------
-------------
2.2.2.1         ethernet1/1/1        v2/S           1            30               1
2.2.2.2
1.1.1.1         ethernet1/1/26:1     v2/S           1            30               1
1.1.1.2
15.1.1.1        vlan2001             v2/S           0            30               1
15.1.1.1
```

The `show ip pim neighbor` command displays the PIM neighbor of LHR and the interface to reach the neighbor.

```
LHR# show ip pim neighbor
Neighbor Address         Interface            Uptime/Expires         Ver        DR Prio/Mode
-------------------------------------------------------------------------------------------
2.2.2.2                  ethernet1/1/17      00:02:58/00:01:24       v2           1   / DR S
1.1.1.2                  ethernet1/1/29      00:07:49/00:01:31       v2           1   / DR S
```

```
LHR# show ip pim rp mapping
Group(s) : 224.0.0.0/4,  Static
RP : 192.168.1.25, v2
```

The following `show` command output examples display the PIM states across all nodes after IGMP join and multicast traffic is received.

PIM states in FHR node

The `show ip pim tib` command output displays the PIM tree information base (TIB).

```
FHR# show ip pim tib

PIM Multicast Routing Table
Flags: S - Sparse, C - Connected, L - Local, P - Pruned,
       R - RP-bit set, F - Register Flag, T - SPT-bit set, J - Join SPT,
       K - Ack-Pending state
Timers: Uptime/Expires
Interface state: Interface, next-Hop, State/Mode

(22.1.1.10, 224.1.1.1), uptime 00:02:20, expires 00:01:09, flags: T
  Incoming interface: ethernet1/1/48, RPF neighbor 0.0.0.0
  Outgoing interface list:
    ethernet1/1/17  Forward/Sparse  00:00:19/00:03:10
```

The `show ip pim mcache` command output displays multicast route entries.

```
FHR# show ip pim mcache
PIM Multicast Routing Cache Table

(22.1.1.10,224.1.1.1)
  Incoming interface : ethernet1/1/48
  Outgoing interface list :
   ethernet1/1/17
```

PIM states in RP node

```
RP# show ip pim tib

PIM Multicast Routing Table
Flags: S - Sparse, C - Connected, L - Local, P - Pruned,
       R - RP-bit set, F - Register Flag, T - SPT-bit set, J - Join SPT,
       K - Ack-Pending state
Timers: Uptime/Expires
Interface state: Interface, next-Hop, State/Mode

(*, 224.1.1.1), uptime 00:00:57, expires 00:00:00, RP 192.168.1.25, flags: S
  Incoming interface:  Null, RPF neighbor 0.0.0.0
  Outgoing interface list:

(22.1.1.10, 224.1.1.1), uptime 00:02:58, expires 00:03:06, flags: P
  Incoming interface: ethernet1/1/31, RPF neighbor 3.3.3.2
  Outgoing interface list:
```

IGMP and PIM states in LHR node

The `show ip igmp groups` command output displays the IGMP database.

```
LHR# show ip igmp groups
Total Number of Groups: 1
IGMP Connected Group Membership
Group Address                 Interface           Mode                Uptime
Expires             Last Reporter
224.1.1.1                     vlan2001            IGMPv2-Compat       00:00:01
00:01:59            15.1.1.10
LHR#
```

```
LHR# show ip pim tib

PIM Multicast Routing Table
Flags: S - Sparse, C - Connected, L - Local, P - Pruned,
       R - RP-bit set, F - Register Flag, T - SPT-bit set, J - Join SPT,
       K - Ack-Pending state
Timers: Uptime/Expires
Interface state: Interface, next-Hop, State/Mode

(*, 224.1.1.1), uptime 00:00:05, expires 00:00:54, RP 192.168.1.25, flags: SCJ
  Incoming interface: ethernet1/1/29, RPF neighbor 1.1.1.2
  Outgoing interface list:
    vlan2001  Forward/Sparse  00:00:05/Never

(22.1.1.10, 224.1.1.1), uptime 00:00:05, expires 00:03:24, flags: T
  Incoming interface: ethernet1/1/17, RPF neighbor 2.2.2.2
  Outgoing interface list:
    vlan2001  Forward/Sparse  00:00:05/Never
```

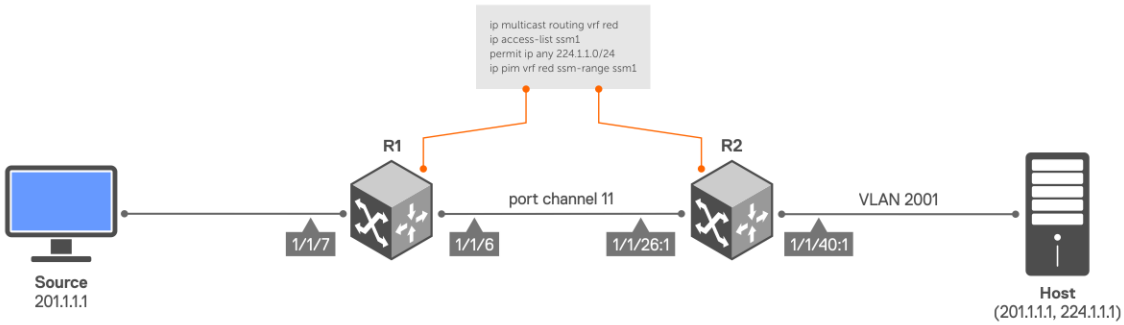
```
LHR# show ip pim mcache
PIM Multicast Routing Cache Table

(*, 224.1.1.1)
  Incoming interface : ethernet1/1/29
  Outgoing interface list :
   vlan2001
```

```
(22.1.1.10,224.1.1.1)
  Incoming interface : ethernet1/1/17
  Outgoing interface list :
   vlan2001
```

PIM-SSM sample configuration

This section describes how to enable PIM-SSM using the topology show in the following illustration.



To enable PIM-SSM, perform the following configurations on R1 and R2:

Sample configuration on R1:

```
R1# configure terminal
R1(config)# ip vrf red
R1(conf-vrf)# end

R1# configure terminal
R1(config)# interface port-channel 11
R1(conf-if-po-11)# no switchport
R1(conf-if-po-11)# ip vrf forwarding red
R1(conf-if-po-11)# end

R1# configure terminal
R1(config)# interface ethernet 1/1/6
R1(conf-if-eth1/1/6)# no ip vrf forwarding
R1(conf-if-eth1/1/6)# no switchport
R1(conf-if-eth1/1/6)# channel-group 11
R1(conf-if-eth1/1/6)# end

R1# configure terminal
R1(config)# interface ethernet 1/1/7
R1(conf-if-eth1/1/7)# no switchport
R1(conf-if-eth1/1/7)# interface ethernet 1/1/7
R1(conf-if-eth1/1/7)# ip vrf forwarding red
R1(conf-if-eth1/1/7)# ip address 201.1.1.2/24
R1(conf-if-eth1/1/7)# ip pim sparse-mode
R1(conf-if-eth1/1/7)# no shutdown
R1(conf-if-eth1/1/7)# end

R1# configure terminal
R1(config)# interface port-channel 11
R1(conf-if-po-11)# no switchport
R1(conf-if-po-11)# interface port-channel 11
R1(conf-if-po-11)# ip vrf forwarding red
R1(conf-if-po-11)# ip address 193.1.1.1/24
R1(conf-if-po-11)# ip pim sparse-mode
R1(conf-if-po-11)# no shutdown
R1(conf-if-po-11)# end

R1# configure terminal
R1(config)# interface Lo0
R1(conf-if-lo-0)# ip vrf forwarding red
```

```
R1(conf-if-lo-0)# ip address 2.2.2.2/32
R1(conf-if-lo-0)# ip pim sparse-mode
R1(conf-if-lo-0)# no shutdown
R1(conf-if-lo-0)# end

R1# configure terminal
R1(config)# router ospf 100 vrf red
R1(config-router-ospf-100)# interface ethernet 1/1/7
R1(conf-if-eth1/1/7)# ip ospf 100 area 0
R1(conf-if-eth1/1/7)# end

R1# configure terminal
R1(config)# router ospf 100 vrf red
R1(config-router-ospf-100)# interface port-channel 11
R1(conf-if-po-11)# ip ospf 100 area 0
R1(conf-if-po-11)# end

R1# configure terminal
R1(config)# ip multicast-routing vrf red
R1(config)# end

R1# configure terminal
R1(config)# ip access-list test
R1(config-ipv4-acl)# permit ip any 224.1.1.0/24
R1(config-ipv4-acl)# exit

R1(config)# ip pim vrf red ssm-range test
R1(config)# end
```

Sample configuration on R2:

```
R2# configure terminal
R2(config)# ip vrf red
R2(conf-vrf)# end

R2# configure terminal
R2(config)# interface vlan 2001
R2(conf-if-vl-2001)# ip vrf forwarding red
R2(conf-if-vl-2001)# end

R2# configure terminal
R2(config)# interface ethernet 1/1/40:1
R2(conf-if-eth1/1/40:1)# no ip vrf forwarding
R2(conf-if-eth1/1/40:1)# switchport mode trunk
R2(conf-if-eth1/1/40:1)# switchport trunk allowed vlan 2001
R2(conf-if-eth1/1/40:1)# end

R2# configure terminal
R2(config)# interface port-channel 11
R2(conf-if-po-11)# no switchport
R2(conf-if-po-11)# ip vrf forwarding red
R2(conf-if-po-11)# end

R2# configure terminal
R2(config)# interface ethernet 1/1/26:1
R2(conf-if-eth1/1/26:1)# no ip vrf forwarding
R2(conf-if-eth1/1/26:1)# no switchport
R2(conf-if-eth1/1/26:1)# channel-group 11
R2(conf-if-eth1/1/26:1)# end

R2# configure terminal
R2(config)# interface vlan 2001
R2(conf-if-vl-2001)# ip vrf forwarding red
R2(conf-if-vl-2001)# ip address 208.1.1.2/24
R2(conf-if-vl-2001)# ip pim sparse-mode
R2(conf-if-vl-2001)# no shutdown
R2(conf-if-vl-2001)# end

R2# configure terminal
```

```
R2(config)# interface port-channel 11
R2(conf-if-po-11)# no switchport
R2(conf-if-po-11)# interface port-channel 11
R2(conf-if-po-11)# ip vrf forwarding red
R2(conf-if-po-11)# ip address 193.1.1.2/24
R2(conf-if-po-11)# ip pim sparse-mode
R2(conf-if-po-11)# no shutdown
R2(conf-if-po-11)# end

R2# configure terminal
R2(config)# interface Lo0
R2(conf-if-lo-0)# ip vrf forwarding red
R2(conf-if-lo-0)# ip address 4.4.4.4/32
R2(conf-if-lo-0)# ip pim sparse-mode
R2(conf-if-lo-0)# no shutdown
R2(conf-if-lo-0)# end

R2# configure terminal
R2(config)# router ospf 100 vrf red
R2(config-router-ospf-100)# interface vlan 2001
R2(conf-if-vl-2001)# ip ospf 100 area 0
R2(conf-if-vl-2001)# end

R2# configure terminal
R2(config)# router ospf 100 vrf red
R2(config-router-ospf-100)# interface port-channel 11
R2(conf-if-po-11)# ip ospf 100 area 0
R2(conf-if-po-11)# end

R2# configure terminal
R2(config)# ip multicast-routing vrf red
R2(config)# end

R2# configure terminal
R2(config)# ip access-list test
R2(config-ipv4-acl)# permit ip any 224.1.1.0/24
R2(config-ipv4-acl)# exit
R2(config)# ip pim vrf red ssm-range test
R2(config)# end
```

Verify the configuration

To verify the configuration, use the following `show` commands on R1:

The `show ip pim vrf red neighbor` command displays the PIM neighbor of R1 and the interface through which the neighbor is reached.

```
R1# show ip pim vrf red neighbor
Neighbor Address         Interface            Uptime/Expires         Ver          DR Priority /
Mode
-------------------------------------------------------------------------------------------------
-------------
193.1.1.2                port-channel11      02:34:33/00:01:17        v2           1          /
DR S
```

The `show ip pim vrf red ssm-range` command displays the specified multicast address range.

```
R1# show ip pim vrf red ssm-range
Group Address   / MaskLen
224.1.1.0       / 24
```

The `show ip pim vrf red tib` command output displays the PIM tree information base (TIB).

```
R1# show ip pim vrf red tib
PIM Multicast Routing Table
Flags: S - Sparse, C - Connected, L - Local, P - Pruned,
       R - RP-bit set, F - Register Flag, T - SPT-bit set, J - Join SPT,
       K - Ack-Pending state
```

```
Timers: Uptime/Expires
Interface state: Interface, next-Hop, State/Mode

(201.1.1.1, 224.1.1.1), uptime 00:19:42, expires 00:00:47, flags: T
  Incoming interface: ethernet1/1/7, RPF neighbor 0.0.0.0
  Outgoing interface list:
    port-channel11  Forward/Sparse  00:00:37/00:02:52
```

The `show ip pim vrf red mcache` command output displays multicast route entries.

```
R1# show ip pim vrf red mcache
PIM Multicast Routing Cache Table
(201.1.1.1, 224.1.1.1)
  Incoming interface : ethernet1/1/7
  Outgoing interface list :
   port-channel11
```

Use the following `show` commands on R2:

The `show ip igmp vrf red groups` command output displays the IGMP database.

```
R2# show ip igmp vrf red groups
Total Number of Groups: 1
IGMP Connected Group Membership
Group Address   Interface   Mode     Uptime    Expires    Last Reporter
224.1.1.1       vlan2001    Include  00:00:03  Never      208.1.1.1
```

The `show ip pim vrf red tib` command output displays the PIM tree information base (TIB).

```
R2# show ip pim vrf red tib
PIM Multicast Routing Table
Flags: S - Sparse, C - Connected, L - Local, P - Pruned,
       R - RP-bit set, F - Register Flag, T - SPT-bit set, J - Join SPT,
       K - Ack-Pending state
Timers: Uptime/Expires
Interface state: Interface, next-Hop, State/Mode

(201.1.1.1, 224.1.1.1), uptime 00:00:06, expires 00:03:23, flags: CT
  Incoming interface: port-channel11, RPF neighbor 193.1.1.1
  Outgoing interface list:
    vlan2001  Forward/Sparse  00:00:06/Never
```

The `show ip pim vrf red neighbor` command displays the PIM neighbor of R2 and the interface through which the neighbor is reached.

```
R2# show ip pim vrf red neighbor
Neighbor Address  Interface      Uptime/Expires    Ver DR Priority / Mode
--------------------------------------------------------------------------
193.1.1.1         port-channel11 02:34:15/00:01:29 v2      1/ S
```

The `show ip pim vrf red ssm-range` command displays the specified multicast address range.

```
R2# show ip pim vrf red ssm-range
Group Address   / MaskLen
224.1.1.0       / 24
```

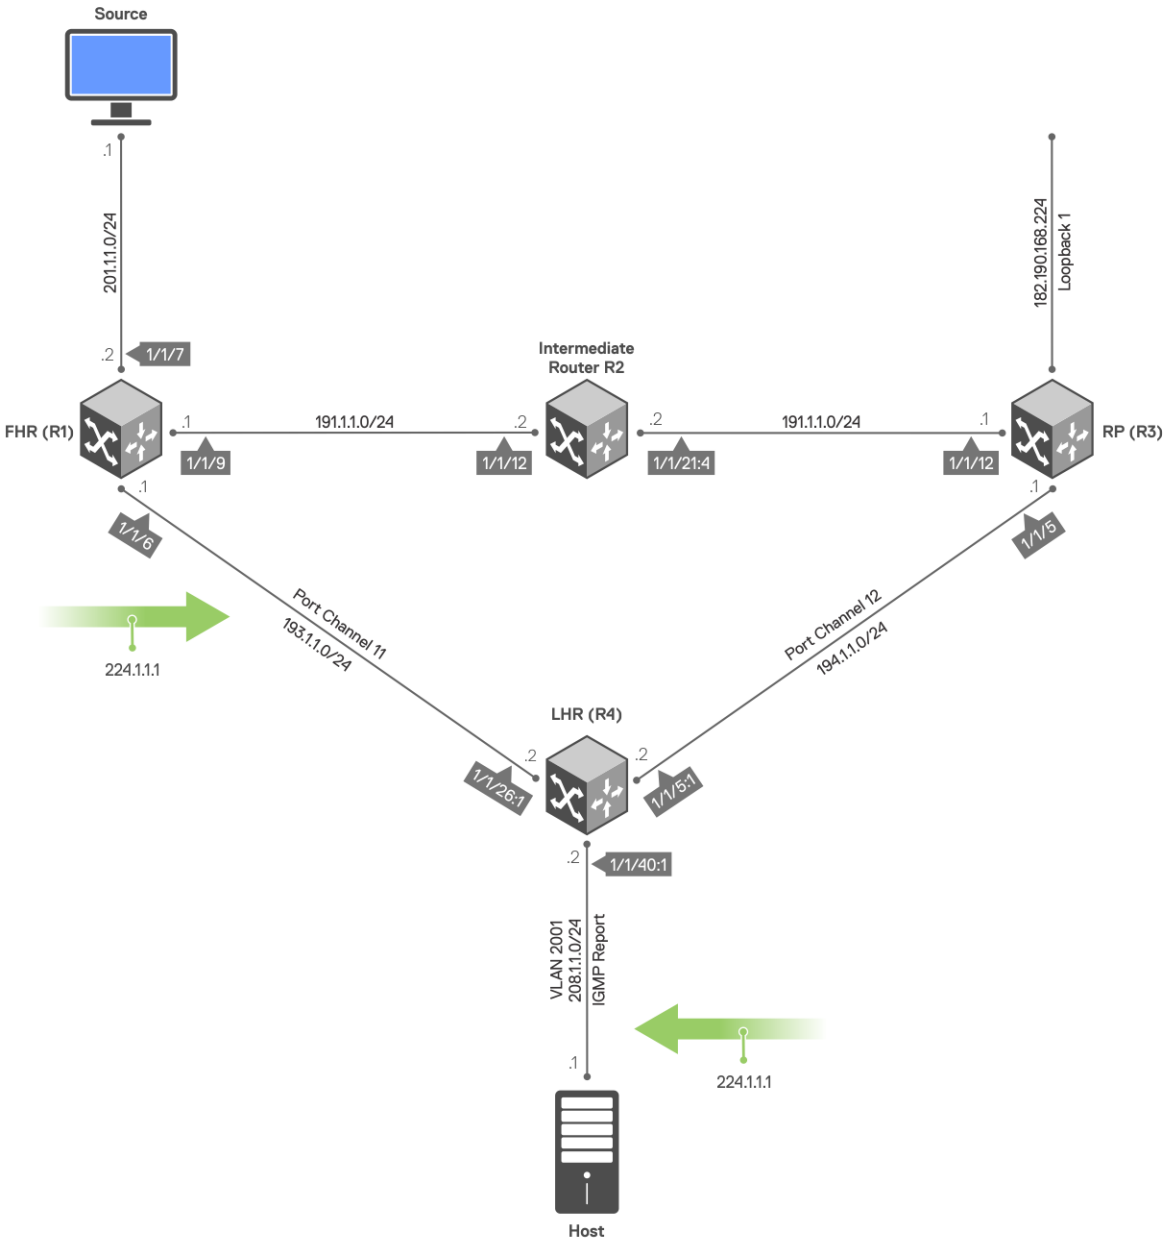
The `show ip pim vrf red mcache` command output displays multicast route entries.

```
R2# show ip pim vrf red mcache
PIM Multicast Routing Cache Table

(201.1.1.1, 224.1.1.1)
  Incoming interface : port-channel11
  Outgoing interface list :
   vlan2001
```

Multicast VRF sample configuration

This section describes how to configure IPv4 multicast in a non-default VRF instance using the topology shown in the following illustration.



Perform the following configuration on each of the nodes, R1, R2, R3, and R4.

Sample configuration on R1:

```
R1# configure terminal
R1(config)# ip vrf red
R1(conf-vrf)# end

R1# configure terminal
R1(config)# interface port-channel 11
R1(conf-if-po-11)# no switchport
```

```
R1(conf-if-po-11)# ip vrf forwarding red
R1(conf-if-po-11)# end

R1# configure terminal
R1(config)# interface ethernet 1/1/6
R1(conf-if-eth1/1/6)# no ip vrf forwarding
R1(conf-if-eth1/1/6)# no switchport
R1(conf-if-eth1/1/6)# channel-group 11
R1(conf-if-eth1/1/6)# end

R1# configure terminal
R1(config)# interface ethernet 1/1/7
R1(conf-if-eth1/1/7)# no switchport
R1(conf-if-eth1/1/7)# interface ethernet 1/1/7
R1(conf-if-eth1/1/7)# ip vrf forwarding red
R1(conf-if-eth1/1/7)# ip address 201.1.1.2/24
R1(conf-if-eth1/1/7)# ip pim sparse-mode
R1(conf-if-eth1/1/7)# no shutdown
R1(conf-if-eth1/1/7)# end

R1# configure terminal
R1(config)# interface ethernet 1/1/9
R1(conf-if-eth1/1/9)# no switchport
R1(conf-if-eth1/1/9)# interface ethernet 1/1/9
R1(conf-if-eth1/1/9)# ip vrf forwarding red
R1(conf-if-eth1/1/9)# ip address 191.1.1.1/24
R1(conf-if-eth1/1/9)# ip pim sparse-mode
R1(conf-if-eth1/1/9)# no shutdown
R1(conf-if-eth1/1/9)# end

R1# configure terminal
R1(config)# interface port-channel 11
R1(conf-if-po-11)# no switchport
R1(conf-if-po-11)# interface port-channel 11
R1(conf-if-po-11)# ip vrf forwarding red
R1(conf-if-po-11)# ip address 193.1.1.1/24
R1(conf-if-po-11)# ip pim sparse-mode
R1(conf-if-po-11)# no shutdown
R1(conf-if-po-11)# end

R1# configure terminal
R1(config)# interface Lo0
R1(conf-if-lo-0)# ip vrf forwarding red
R1(conf-if-lo-0)# ip address 2.2.2.2/32
R1(conf-if-lo-0)# ip pim sparse-mode
R1(conf-if-lo-0)# no shutdown
R1(conf-if-lo-0)# end

R1# configure terminal
R1(config)# router ospf 100 vrf red
R1(config-router-ospf-100)# interface ethernet 1/1/7
R1(conf-if-eth1/1/7)# ip ospf 100 area 0
R1(conf-if-eth1/1/7)# end

R1# configure terminal
R1(config)# router ospf 100 vrf red
R1(config-router-ospf-100)# interface ethernet 1/1/9
R1(conf-if-eth1/1/9)# ip ospf 100 area 0
R1(conf-if-eth1/1/9)# end

R1# configure terminal
R1(config)# router ospf 100 vrf red
R1(config-router-ospf-100)# interface port-channel 11
R1(conf-if-po-11)# ip ospf 100 area 0
R1(conf-if-po-11)# end

R1# configure terminal
R1(config)# ip multicast-routing vrf red
R1(config)# end
```

```
R1# configure terminal
R1(config)# ip pim vrf red rp-address 182.190.168.224 group-address 224.0.0.0/4
R1(config)# end
```

Sample configuration on R2:

```
R2# configure terminal
R2(config)# ip vrf red
R2(conf-vrf)# end

R2# configure terminal
R2(config)# interface vlan 1001
R2(conf-if-vl-1001)# ip vrf forwarding red
R2(conf-if-vl-1001)# end

R2# configure terminal
R2(config)# interface ethernet 1/1/21:4
R2(conf-if-eth1/1/21:4)# switchport mode trunk
R2(conf-if-eth1/1/21:4)# switchport trunk allowed vlan 1001
R2(conf-if-eth1/1/21:4)# end

R2# configure terminal
R2(config)# interface ethernet 1/1/12:1
R2(conf-if-eth1/1/12:1)# no switchport
R2(conf-if-eth1/1/12:1)# ip vrf forwarding red
R2(conf-if-eth1/1/12:1)# ip address 191.1.1.2/24
R2(conf-if-eth1/1/12:1)# ip pim sparse-mode
R2(conf-if-eth1/1/12:1)# no shutdown
R2(conf-if-eth1/1/12:1)# end

R2# configure terminal
R2(config)# interface vlan 1001
R2(conf-if-vl-1001)# ip vrf forwarding red
R2(conf-if-vl-1001)# ip address 192.1.1.2/24
R2(conf-if-vl-1001)# ip pim sparse-mode
R2(conf-if-vl-1001)# no shutdown
R2(conf-if-vl-1001)# end

R2# configure terminal
R2(config)# interface Lo0
R2(conf-if-lo-0)# ip vrf forwarding red
R2(conf-if-lo-0)# ip address 1.1.1.1/32
R2(conf-if-lo-0)# ip pim sparse-mode
R2(conf-if-lo-0)# no shutdown
R2(conf-if-lo-0)# end

R2# configure terminal
R2(config)# router ospf 100 vrf red
R2(config-router-ospf-100)# interface ethernet 1/1/12:1
R2(conf-if-eth1/1/12:1)# ip ospf 100 area 0
R2(conf-if-eth1/1/12:1)# end

R2# configure terminal
R2(config)# router ospf 100 vrf red
R2(config-router-ospf-100)# interface vlan 1001
R2(conf-if-vl-1001)# ip ospf 100 area 0
R2(conf-if-vl-1001)# end

R2# configure terminal
R2(config)# ip multicast-routing vrf red
R2(config)# end

R2# configure terminal
R2(config)# ip pim vrf red rp-address 182.190.168.224 group-address 224.0.0.0/4
R2(config)# end
```

Sample configuration on R3:

```
R3# configure terminal
R3(config)# ip vrf red
R3(conf-vrf)# end

R3# configure terminal
R3(config)# interface vlan 1001
R3(conf-if-vl-1001)# ip vrf forwarding red
R3(conf-if-vl-1001)# end

R3# configure terminal
R3(config)# interface ethernet 1/1/12
R3(conf-if-eth1/1/12)# no ip vrf forwarding
R3(conf-if-eth1/1/12)# switchport mode trunk
R3(conf-if-eth1/1/12)# switchport trunk allowed vlan 1001
R3(conf-if-eth1/1/12)# end

R3# configure terminal
R3(config)# interface port-channel 12
R3(conf-if-po-12)# no switchport
R3(conf-if-po-12)# ip vrf forwarding red
R3(conf-if-po-12)# end
R3# configure terminal
R3(config)# interface ethernet 1/1/5
R3(conf-if-eth1/1/5)# no ip vrf forwarding
R3(conf-if-eth1/1/5)# no switchport
R3(conf-if-eth1/1/5)# channel-group 12
R3(conf-if-eth1/1/5)# end

R3# configure terminal
R3(config)# interface vlan 1001
R3(conf-if-vl-1001)# ip vrf forwarding red
R3(conf-if-vl-1001)# ip address 192.1.1.1/24
R3(conf-if-vl-1001)# ip pim sparse-mode
R3(conf-if-vl-1001)# no shutdown
R3(conf-if-vl-1001)# end

R3# configure terminal
R3(config)# interface port-channel 12
R3(conf-if-po-12)# no switchport
R3(conf-if-po-12)# interface port-channel 12
R3(conf-if-po-12)# ip vrf forwarding red
R3(conf-if-po-12)# ip address 194.1.1.1/24
R3(conf-if-po-12)# ip pim sparse-mode
R3(conf-if-po-12)# no shutdown
R3(conf-if-po-12)# end

R3# configure terminal
R3(config)# interface Lo0
R3(conf-if-lo-0)# ip vrf forwarding red
R3(conf-if-lo-0)# ip address 3.3.3.3/32
R3(conf-if-lo-0)# ip pim sparse-mode
R3(conf-if-lo-0)# no shutdown
R3(conf-if-lo-0)# end

R3# configure terminal
R3(config)# router ospf 100 vrf red
R3(config-router-ospf-100)# interface vlan 1001
R3(conf-if-vl-1001)# ip ospf 100 area 0
R3(conf-if-vl-1001)# end

R3# configure terminal
R3(config)# router ospf 100 vrf red
R3(config-router-ospf-100)# interface port-channel 12
R3(conf-if-po-12)# ip ospf 100 area 0
R3(conf-if-po-12)# end

R3# configure terminal
```

```
R3(config)# router ospf 100 vrf red
R3(config-router-ospf-100)# interface Lo1
R3(conf-if-lo-1)# ip ospf 100 area 0
R3(conf-if-lo-1)# end

R3# configure terminal
R3(config)# ip multicast-routing vrf red
R3(config)# end

R3# configure terminal
R3(config)# interface Lo1
R3(conf-if-lo-1)# ip vrf forwarding red
R3(conf-if-lo-1)# ip address 182.190.168.224/32
R3(conf-if-lo-1)# ip pim sparse-mode
R3(conf-if-lo-1)# no shutdown
R3(conf-if-lo-1)# end

R3# configure terminal
R3(config)# ip pim vrf red rp-address 182.190.168.224 group-address 224.0.0.0/4
R3(config)# end
```

Sample configuration on R4:

```
R4# configure terminal
R4(config)# ip vrf red
R4(conf-vrf)# end

R4# configure terminal
R4(config)# interface vlan 2001
R4(conf-if-vl-2001)# ip vrf forwarding red
R4(conf-if-vl-2001)# end

R4# configure terminal
R4(config)# interface ethernet 1/1/40:1
R4(conf-if-eth1/1/40:1)# no ip vrf forwarding
R4(conf-if-eth1/1/40:1)# switchport mode trunk
R4(conf-if-eth1/1/40:1)# switchport trunk allowed vlan 2001
R4(conf-if-eth1/1/40:1)# end

R4# configure terminal
R4(config)# interface port-channel 11
R4(conf-if-po-11)# no switchport
R4(conf-if-po-11)# ip vrf forwarding red
R4(conf-if-po-11)# end
R4# configure terminal
R4(config)# interface port-channel 12
R4(conf-if-po-12)# no switchport
R4(conf-if-po-12)# ip vrf forwarding red
R4(conf-if-po-12)# end

R4# configure terminal
R4(config)# interface ethernet 1/1/26:1
R4(conf-if-eth1/1/26:1)# no ip vrf forwarding
R4(conf-if-eth1/1/26:1)# no switchport
R4(conf-if-eth1/1/26:1)# channel-group 11
R4(conf-if-eth1/1/26:1)# end

R4# configure terminal
R4(config)# interface ethernet 1/1/5:1
R4(conf-if-eth1/1/5:1)# no ip vrf forwarding
R4(conf-if-eth1/1/5:1)# no switchport
R4(conf-if-eth1/1/5:1)# channel-group 12
R4(conf-if-eth1/1/5:1)# end

R4# configure terminal
R4(config)# interface vlan 2001
R4(conf-if-vl-2001)# ip vrf forwarding red
R4(conf-if-vl-2001)# ip address 208.1.1.2/24
R4(conf-if-vl-2001)# ip pim sparse-mode
```

```
R4(conf-if-vl-2001)# no shutdown
R4(conf-if-vl-2001)# end


R4# configure terminal
R4(config)# interface port-channel 11
R4(conf-if-po-11)# no switchport
R4(conf-if-po-11)# interface port-channel 11
R4(conf-if-po-11)# ip vrf forwarding red
R4(conf-if-po-11)# ip address 193.1.1.2/24
R4(conf-if-po-11)# ip pim sparse-mode
R4(conf-if-po-11)# no shutdown
R4(conf-if-po-11)# end

R4# configure terminal
R4(config)# interface port-channel 12
R4(conf-if-po-12)# no switchport
R4(conf-if-po-12)# interface port-channel 12
R4(conf-if-po-12)# ip vrf forwarding red
R4(conf-if-po-12)# ip address 194.1.1.2/24
R4(conf-if-po-12)# ip pim sparse-mode
R4(conf-if-po-12)# no shutdown
R4(conf-if-po-12)# end

R4# configure terminal
R4(config)# interface Lo0
R4(conf-if-lo-0)# ip vrf forwarding red
R4(conf-if-lo-0)# ip address 4.4.4.4/32
R4(conf-if-lo-0)# ip pim sparse-mode
R4(conf-if-lo-0)# no shutdown
R4(conf-if-lo-0)# end


R4# configure terminal
R4(config)# router ospf 100 vrf red
R4(config-router-ospf-100)# interface vlan 2001
R4(conf-if-vl-2001)# ip ospf 100 area 0
R4(conf-if-vl-2001)# end

R4# configure terminal
R4(config)# router ospf 100 vrf red
R4(config-router-ospf-100)# interface port-channel 11
R4(conf-if-po-11)# ip ospf 100 area 0
R4(conf-if-po-11)# end

R4# configure terminal
R4(config)# router ospf 100 vrf red
R4(config-router-ospf-100)# interface port-channel 12
R4(conf-if-po-12)# ip ospf 100 area 0
R4(conf-if-po-12)# end

R4# configure terminal
R4(config)# ip multicast-routing vrf red
R4(config)# end

R4# configure terminal
R4(config)# ip pim vrf red rp-address 182.190.168.224 group-address 224.0.0.0/4
R4(config)# end
```

Verify the configuration

To verify the configuration, use the following `show` commands.

First hop router (R1)

```
R1# show ip pim vrf red neighbor
Neighbor Address  Interface       Uptime/Expires    Ver  DR Priority / Mode
--------------------------------------------------------------------------
```

```
191.1.1.2         ethernet1/1/9   02:13:21/00:01:25 v2      1/ DR S
193.1.1.2         port-channel11  02:15:29/00:01:22 v2      1/ DR S
```

```
R1# show ip pim vrf red tib

PIM Multicast Routing Table
Flags: S - Sparse, C - Connected, L - Local, P - Pruned,
       R - RP-bit set, F - Register Flag, T - SPT-bit set, J - Join SPT,
       K - Ack-Pending state
Timers: Uptime/Expires
Interface state: Interface, next-Hop, State/Mode

(201.1.1.1, 224.1.1.1), uptime 00:00:33, expires 00:02:56, flags: FT
  Incoming interface: ethernet1/1/7, RPF neighbor 0.0.0.0
  Outgoing interface list:
    port-channel11  Forward/Sparse  00:00:33/00:02:56
```

```
R1# show ip pim vrf red rp
Group           RP
--------------------------------
224.1.1.1       182.190.168.224
```

```
R1# show ip pim vrf red rp mapping
Group(s) : 224.0.0.0/4,  Static
RP : 182.190.168.224, v2
```

```
R1# show ip pim vrf red mcache
PIM Multicast Routing Cache Table

(201.1.1.1, 224.1.1.1)
  Incoming interface : ethernet1/1/7
  Outgoing interface list :
   port-channel11
```

Rendezvous point (R3)

```
R3# show ip pim vrf red neighbor
Neighbor Address Interface       Uptime/Expires      Ver  DR Priority / Mode
-------------------------------------------------------------------------------
192.1.1.2        vlan1001        02:11:46/00:01:33   v2      1/ DR S
194.1.1.2        port-channel12  02:14:12/00:01:33   v2      1/ DR S
```

```
R3# show ip pim vrf red tib

PIM Multicast Routing Table
Flags: S - Sparse, C - Connected, L - Local, P - Pruned,
       R - RP-bit set, F - Register Flag, T - SPT-bit set, J - Join SPT,
       K - Ack-Pending state
Timers: Uptime/Expires
Interface state: Interface, next-Hop, State/Mode

(*, 224.1.1.1), uptime 00:01:48, expires 00:00:00, RP 182.190.168.224, flags: S
  Incoming interface:  Null, RPF neighbor 0.0.0.0
  Outgoing interface list:
    port-channel12  Forward/Sparse  00:01:48/00:02:41
```

```
R3# show ip pim vrf red mcache
PIM Multicast Routing Cache Table

(*, 224.1.1.1)
  Incoming interface :
  Outgoing interface list :
   port-channel12
```

```
R3# show ip rpf vrf red
RPF information for 182.190.168.224
   RPF interface:
   RPF neighbor: 0.0.0.0
```

```
   RPF route/mask: 0.0.0.0/0.0.0.0
   RPF type: Unicast
```

```
R3# show ip pim vrf red rp mapping
Group(s) : 224.0.0.0/4,  Static
RP : 182.190.168.224, v2
```

```
R3# show ip pim vrf red rp
Group           RP
---------------------------------
224.1.1.1       182.190.168.224
```

```
R3# show ip pim vrf red rp
Group           RP
---------------------------------
224.1.1.1       182.190.168.224
```

```
R3# show ip pim vrf red tib

PIM Multicast Routing Table
Flags: S - Sparse, C - Connected, L - Local, P - Pruned,
       R - RP-bit set, F - Register Flag, T - SPT-bit set, J - Join SPT,
       K - Ack-Pending state
Timers: Uptime/Expires
Interface state: Interface, next-Hop, State/Mode

(*, 224.1.1.1), uptime 00:04:41, expires 00:00:00, RP 182.190.168.224, flags: S
  Incoming interface:  Null, RPF neighbor 0.0.0.0
  Outgoing interface list:
    port-channel12  Forward/Sparse  00:04:41/00:02:48


(201.1.1.1, 224.1.1.1), uptime 00:01:55, expires 00:02:29, flags: P
  Incoming interface: port-channel12, RPF neighbor 194.1.1.2
  Outgoing interface list:
```

```
R3# show ip pim vrf red mcache
PIM Multicast Routing Cache Table

(*, 224.1.1.1)
  Incoming interface :
  Outgoing interface list :
   port-channel12
```

Last hop router (R4)

```
R4# show ip pim vrf red neighbor
Neighbor Address         Interface            Uptime/Expires         Ver          DR Priority /
Mode
-------------------------------------------------------------------------------------------------
-------------
193.1.1.1                port-channel11      02:11:48/00:01:26        v2           1          /
S
194.1.1.1                port-channel12      02:12:07/00:01:41        v2           1          /
S

R4# show ip pim vrf red rp mapping
Group(s) : 224.0.0.0/4,  Static
RP : 182.190.168.224, v2

R4# show ip pim vrf red rp
Group           RP
---------------------------------
224.1.1.1       182.190.168.224

R4# show ip igmp vrf red groups
Total Number of Groups: 1
IGMP Connected Group Membership
Group Address                Interface           Mode                Uptime
Expires             Last Reporter
```

```
224.1.1.1                       vlan2001            IGMPv2-Compat       00:00:18
00:02:07           208.1.1.1

R4# show ip rpf vrf red
RPF information for 182.190.168.224
   RPF interface: port-channel12
   RPF neighbor: 194.1.1.1
   RPF route/mask: 182.190.168.224/255.255.255.255
   RPF type: Unicast

R4# show ip pim vrf red tib

PIM Multicast Routing Table
Flags: S - Sparse, C - Connected, L - Local, P - Pruned,
       R - RP-bit set, F - Register Flag, T - SPT-bit set, J - Join SPT,
       K - Ack-Pending state
Timers: Uptime/Expires
Interface state: Interface, next-Hop, State/Mode

(*, 224.1.1.1), uptime 00:05:44, expires 00:00:15, RP 182.190.168.224, flags: SCJ
  Incoming interface: port-channel12, RPF neighbor 194.1.1.1
  Outgoing interface list:
    vlan2001  Forward/Sparse  00:05:44/Never


(201.1.1.1, 224.1.1.1), uptime 00:02:58, expires 00:00:31, flags: CT
  Incoming interface: port-channel11, RPF neighbor 193.1.1.1
  Outgoing interface list:
    vlan2001  Forward/Sparse  00:02:58/Never

R4# show ip pim vrf red mcache
PIM Multicast Routing Cache Table

(*, 224.1.1.1)
  Incoming interface : port-channel12
  Outgoing interface list :
   vlan2001

(201.1.1.1, 224.1.1.1)
  Incoming interface : port-channel11
  Outgoing interface list :
   vlan2001

R4# show ip pim vrf red summary

Entries in PIM-TIB/MFC: 3/2

Active Modes:
    PIM-SM

Interface summary:
    4 active PIM interfaces
    0 passive PIM interfaces
    2 active PIM neighbor

TIB Summary:
    1/1 (*,G) entries in PIM-TIB/MFC
    1/1 (S,G) entries in PIM-TIB/MFC
    1/0 (S,G,Rpt) entries in PIM-TIB/MFC

    1 RP
    1 sources

Message Summary:
    81268/13033 Joins/Prunes sent/received
    0 Null Register messages received
    0/0 Register-stop messages sent/received

Data path event summary:
```

```
    0/0 pim-assert messages sent/received
    0/0 register messages sent/received
```

Multicast support on VLT

OS10 supports multicast in a VLT domain in active-standby mode.

In a VLT domain that is in active-standby mode, the designated router (DR) routes multicast traffic (Layer 3) and the other peer VLT node switches (Layer 2) incoming multicast traffic over VLTi links.

In last hop router (LHR), the DR is responsible for triggering upstream PIM joins, but the traffic from RP can reach any one of the VLT peers. For example, if the traffic reaches a peer VLT node that has not triggered the upstream join, it switches the incoming multicast traffic over VLTi links which will reach the originating LHR-DR node. The traffic then gets routed at this node. The L3 routed traffic gets replicated over VLTi links for all the receiver VLANs. This traffic is processed by non-DR LHR node and it forwards it to appropriate receivers based on VLT L2 traffic forwarding rules.

The behavior is the same with first hop router (FHR) as well, where PIM registration is initiated by DR. The PIM join could be sent to the DR or the non-DR FHR node. If the incoming multicast traffic reaches the node which receives the PIM join, it routes the packet towards the RP. If the incoming multicast traffic reaches the peer VLT node that has not received the PIM join, it switches the packet over the VLTi link to the other node, which will route the packet towards the RP.

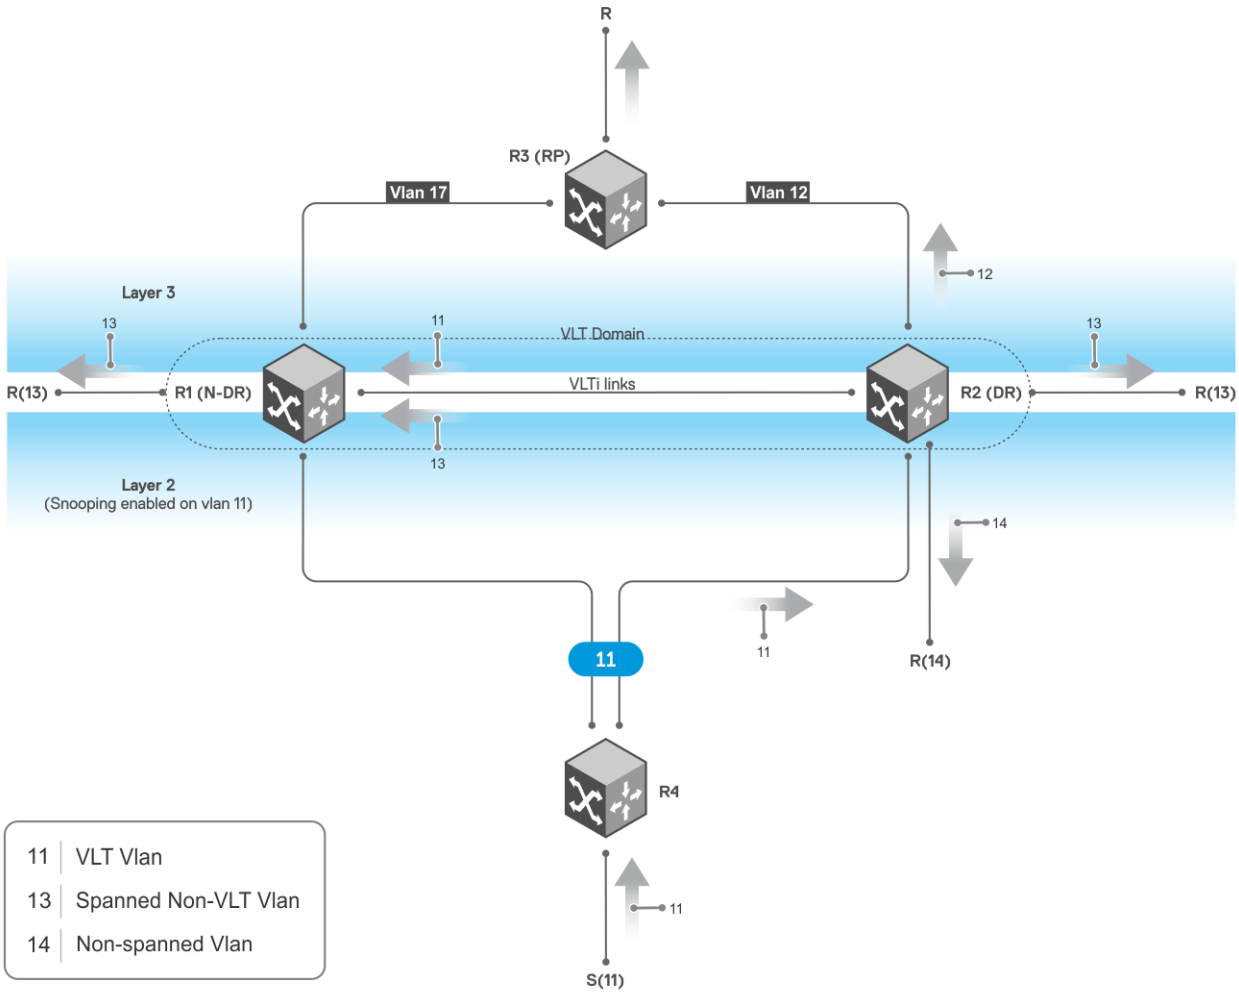
The active-standby mechanism can potentially oversubscribe VLTi links, as multiple copies get replicated for the outgoing receiver VLANs. Dell EMC recommends using only one PIM upstream VLAN towards the RP/source from FHR/LHR.

The following section provides a high-level description of traffic flows for different use cases. The term, spanned VLAN used in the following illustration refers to a VLAN that is configured on both the VLT nodes.

NOTE: Dell EMC recommends that you align the designated router of all spanned VLANs, VLT VLANs and non-VLT VLANs, on to a single VLT node.

Source on VLT VLAN—traffic forwarded to DR

In the following illustration, the source is in VLT VLAN (VLAN 11) and traffic is forwarded to the designated router (R2).



Traffic flow:

1 R4: Traffic from source is switched to VLT LAG, through VLAN 11, and arrives at R2 which is the designated router (DR).

2 R2:

a The (S, G) entry is created.

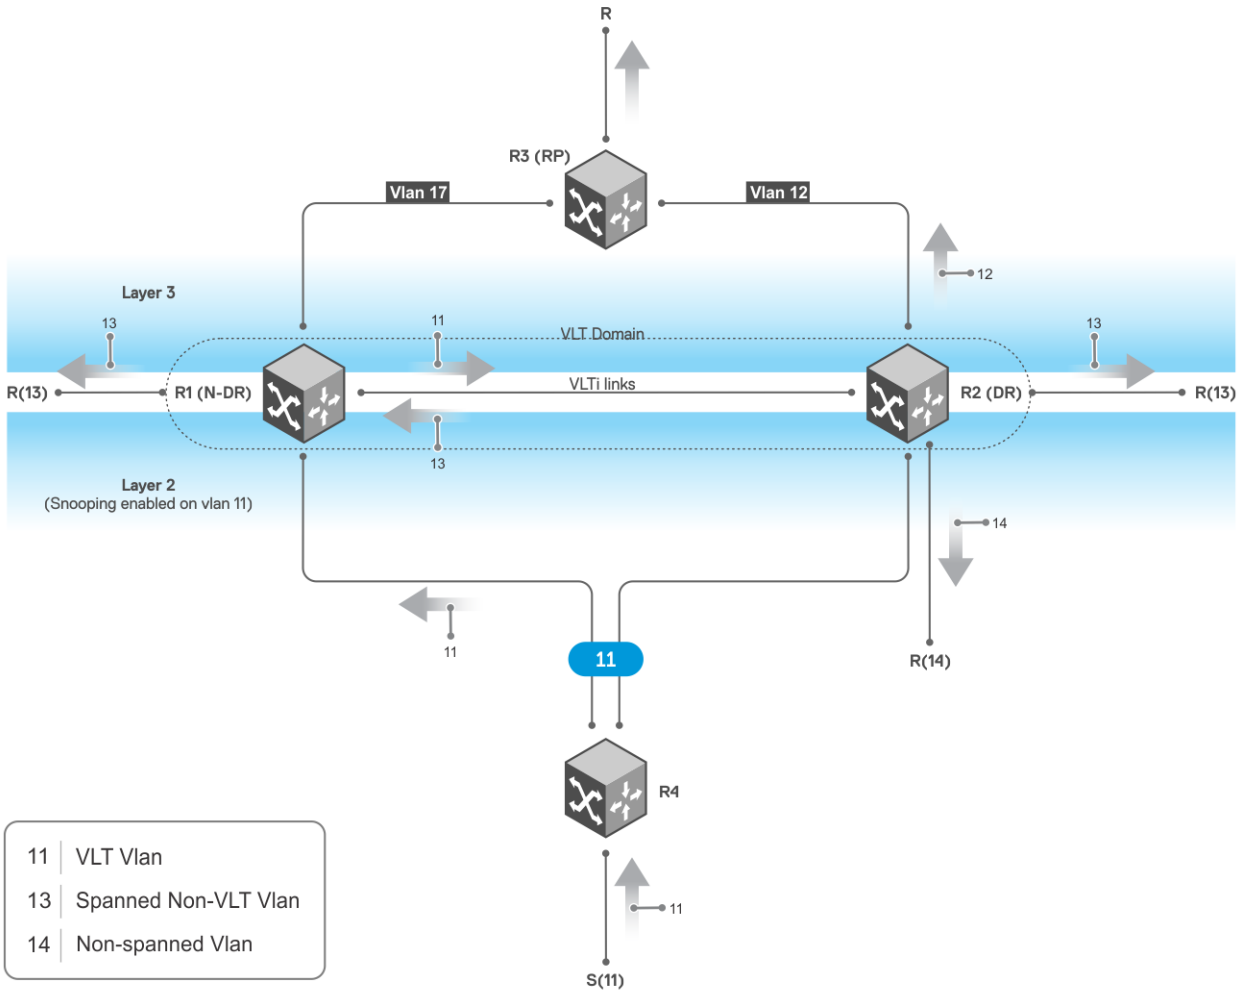
b Traffic is routed to VLAN 12, VLAN 13, and VLAN 14 for receivers R12, R13, and R14.

c Traffic is routed to ICL through VLAN 13, and switched to ICL through VLAN 11.

3 R1: Traffic floods on VLAN 13.

Source on VLT VLAN—traffic forwarded to nonDR

In the following illustration, the source is in VLT VLAN (VLAN 11) and traffic is forwarded to the non-designated router (R1).



Traffic flow:

1 R4: Traffic from source is switched to VLT LAG towards the non-designated router (R1).

2 R1: Traffic is switched to ICL through VLAN 11.

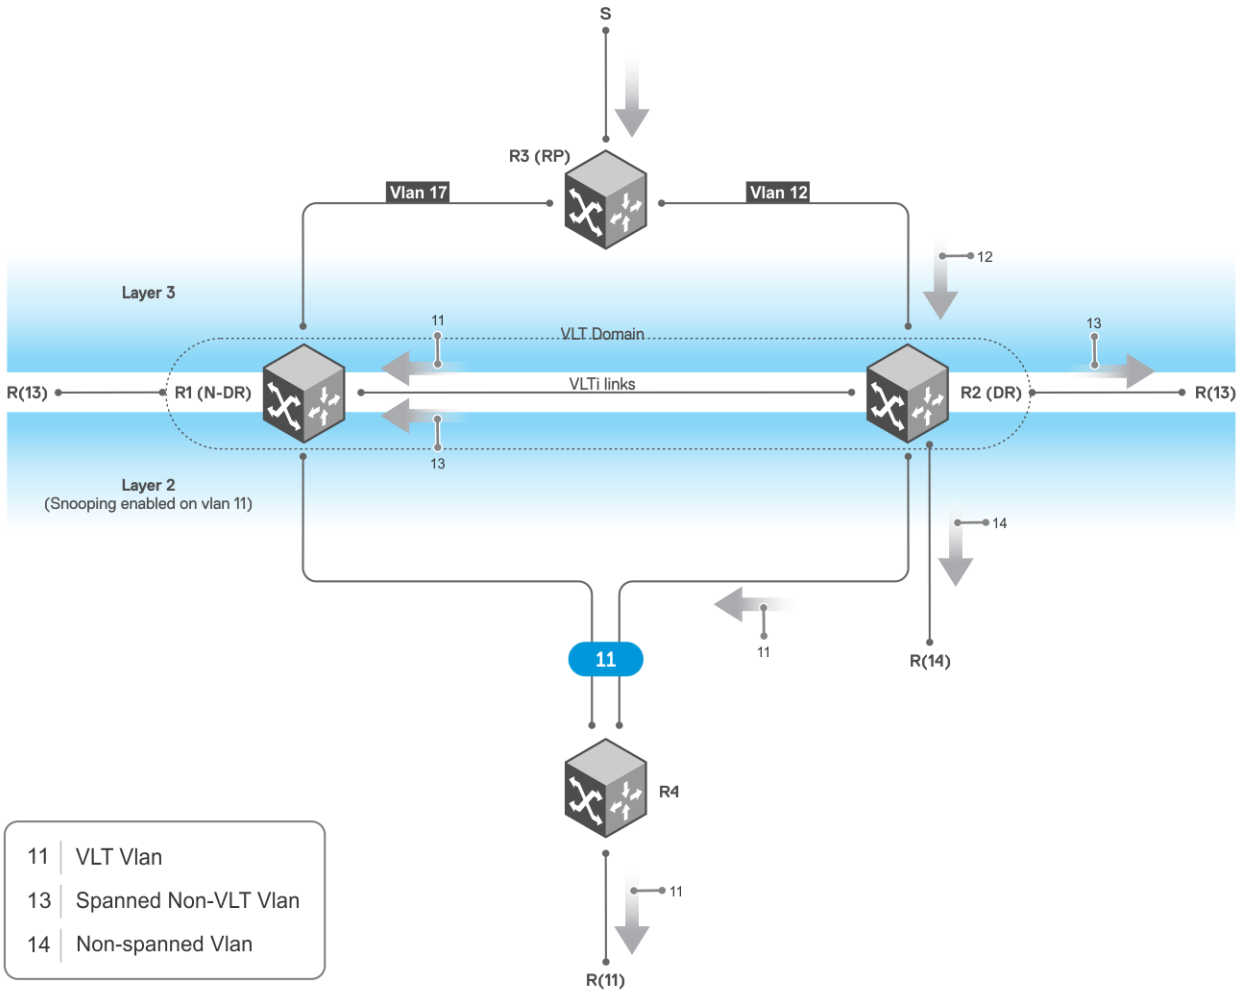
3 R2:

- a The (S, G) entry is created.
- b Traffic is routed to VLAN 12, VLAN 13, and VLAN 14.
- c Traffic is routed to ICL through VLAN 13.

4 Traffic floods on VLAN 13.

Source connected to RP

In the following illustration, the source is connected to the rendezvous point (RP).



Traffic flow:

1 R3: Traffic from source is routed to R2.

2 R2:

a The (S1, G) entry is created.

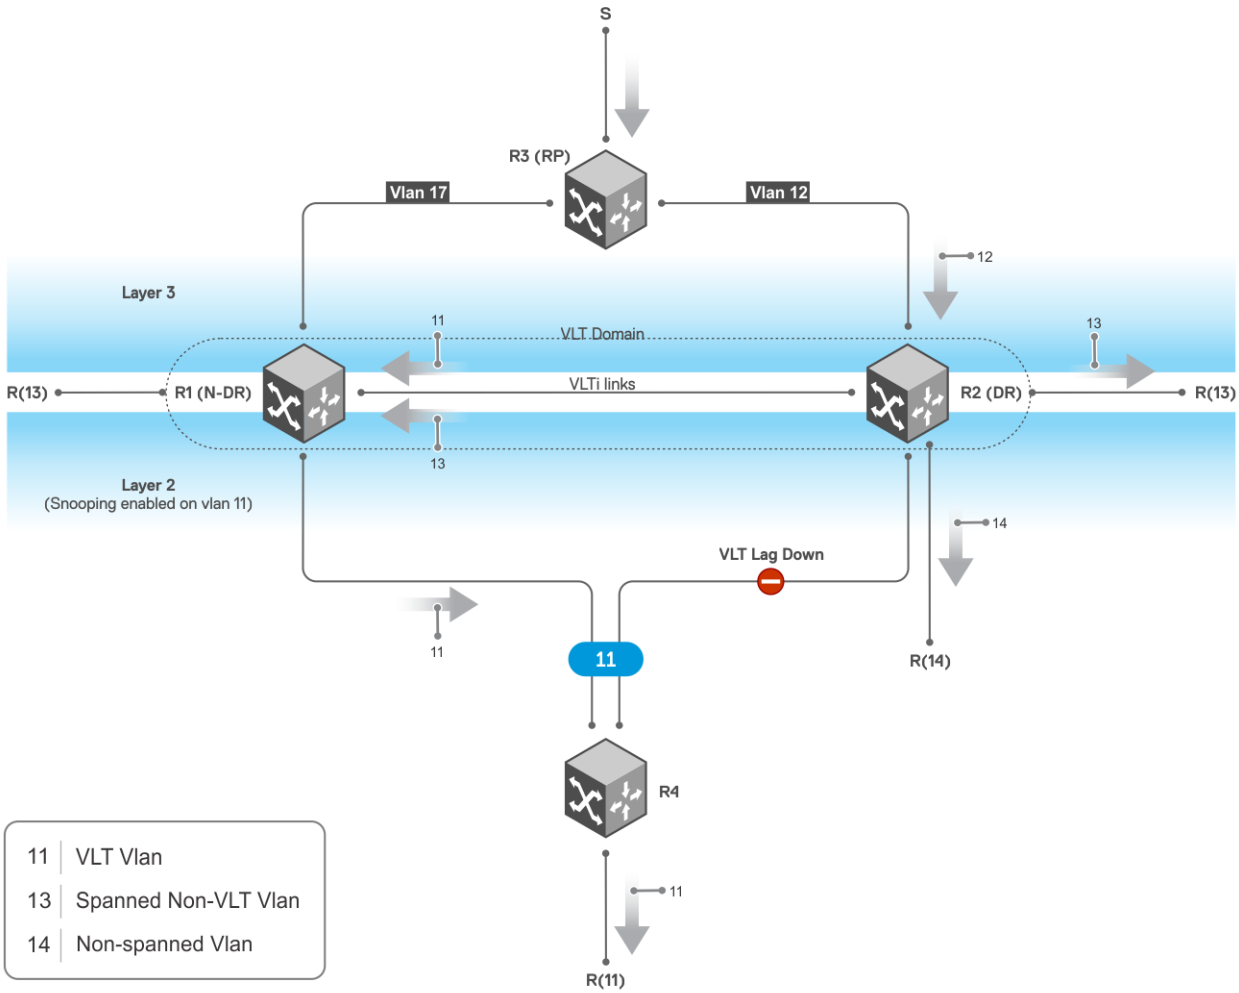
b Traffic is routed to VLAN 11, VLAN 13, and VLAN 14.

c Traffic is routed to ICL through VLAN 11 as well as VLAN 13.

3 Traffic floods on VLAN 13.

VLT LAG down on one side

In the following illustration, VLT LAG is down on one side.



Traffic flow:

1 R3: Traffic from source is routed to R2.

2 R2:

a The (S1, G) entry is created.

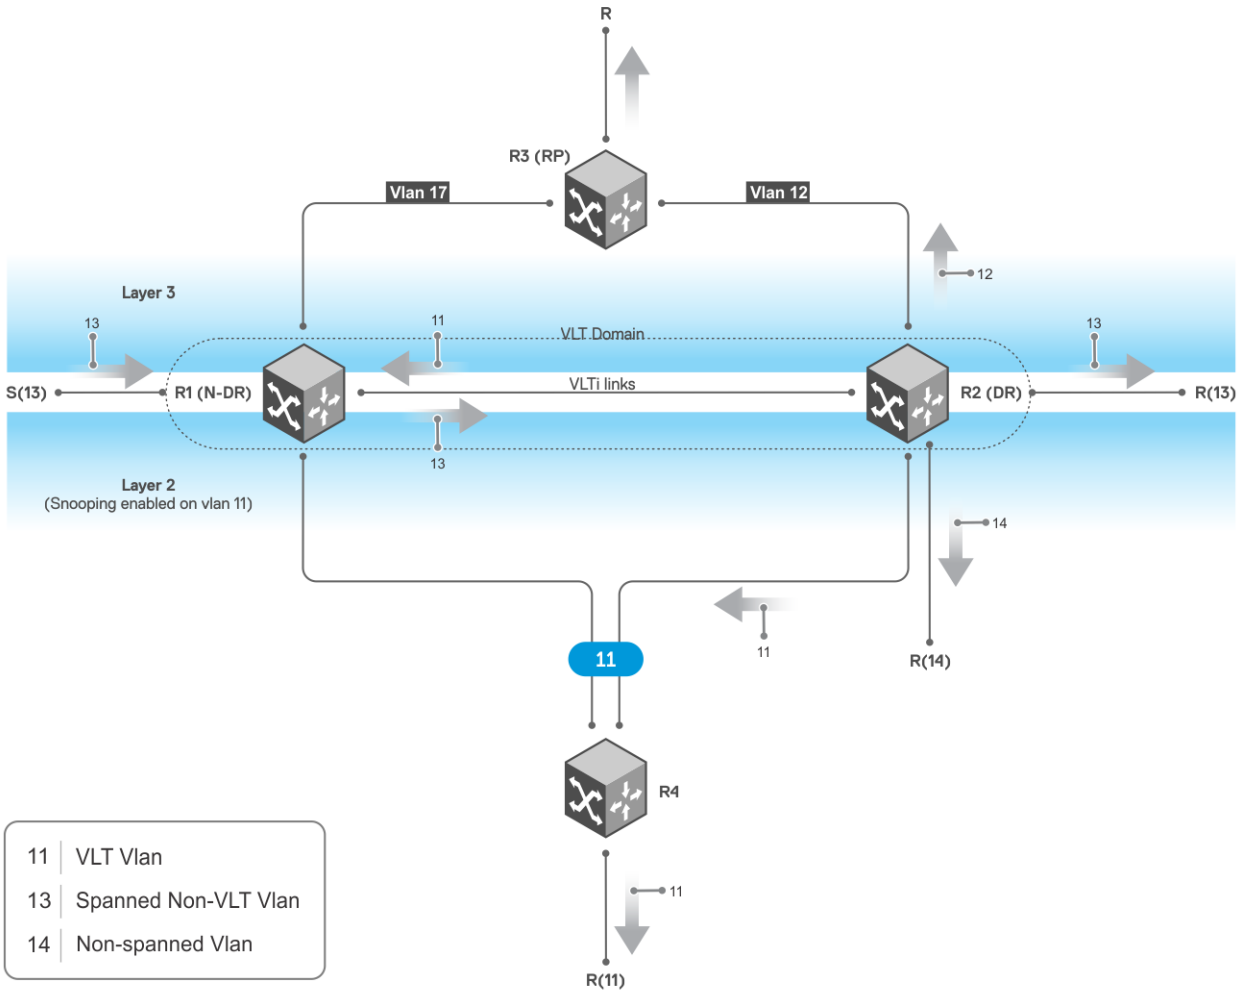
b Traffic is routed to VLAN 11, VLAN 13, and VLAN 14.

c Traffic is routed to ICL through VLAN 11 as well as VLAN 13.

3 Traffic floods on VLAN 13.

Source on spanned non-VLT VLAN

In the following illustration, the source is connected to a router in a spanned non-VLT VLAN.



Traffic flow:

1 R1: Traffic floods to ICL through VLAN 13.

2 R2:

a The (S1, G) entry is created.

b Traffic is routed to VLAN 11, VLAN 12, and VLAN 14.

c Traffic is routed to ICL through VLAN 11.

7

VXLAN

A virtual extensible LAN (VXLAN) extends Layer 2 server connectivity over an underlying Layer 3 transport network in a virtualized data center. A virtualized data center consists of virtual machines (VMs) in a multi-tenant environment. OS10 supports VXLAN as described in RFC 7348.

VXLAN provides a Layer 2 overlay mechanism on an existing Layer 3 network by encapsulating the Layer 2 frames in Layer 3 packets. The VXLAN shared forwarding domain allows hosts (virtual and physical machines) in tenant Layer 2 segments to communicate over the shared IP network. Each tenant Layer 2 segment is identified by a 24-bit ID called a VXLAN network identifier (VNI).

Deployed as a VXLAN gateway, an OS10 switch performs encapsulation/de-encapsulation of L2 frames in L3 packets while tunneling server traffic. In this role, an OS10 switch operates as a VXLAN tunnel endpoint (VTEP). Using VXLAN tunnels, server VLAN segments communicate through the extended L2 forwarding domain.

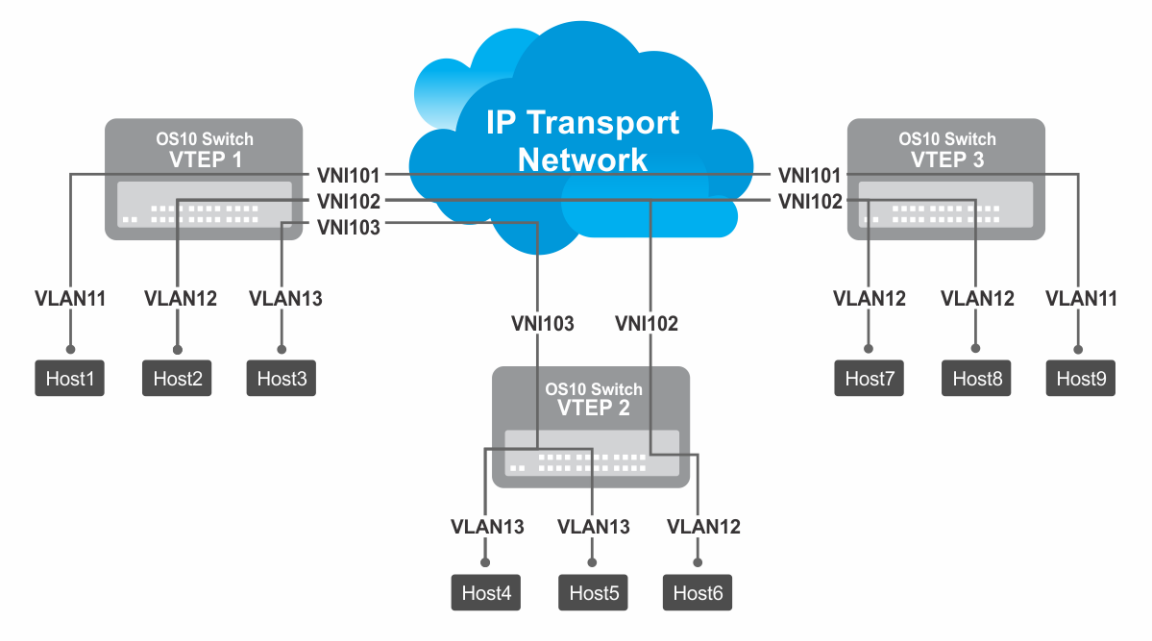


Figure 5. VXLAN topology

VXLAN concepts

Network virtualization overlay (NVO)

An overlay network extends L2 connectivity between server virtual machines (VMs) in a tenant segment over an underlay L3 IP network. A tenant segment can be a group of hosts or servers that are spread across an underlay network.

- The NVO overlay network uses a separate L2 bridge domain (virtual network), which is independent of legacy VLAN forwarding.
- The NVO underlay network operates in the default VRF using the existing L3 infrastructure and routing protocols.

Virtual extensible LAN (VXLAN)	A type of network virtualization overlay that encapsulates a tenant's payload into IP UDP packets for transport across the IP underlay network.
VXLAN network identifier (VNI)	A 24-bit ID number that identifies a tenant segment and is transmitted in a VXLAN encapsulated packet.
VXLAN tunnel endpoint (VTEP)	A switch with connected end hosts that are assigned to virtual networks, and the virtual networks are mapped to VXLAN segments. Local and remote VTEPs performs encapsulation and de-capsulation of VXLAN headers for the traffic between end hosts. A VTEP is sometimes referred to as a network virtualization edge (NVE) node.
Bridge domain	A Layer 2 domain that receives packets from member interfaces and forwards or floods them to other member interfaces based on a packet's destination MAC address. OS10 supports two types of bridge domains: simple VLAN and virtual network. • **Simple VLAN**: A bridge domain represented by a VLAN ID. Traffic on all member ports are assigned with the same VLAN ID. • **Virtual network**: A bridge domain represented by a virtual network ID (VNID). A virtual network supports overlay encapsulation and is mapped with either a single VLAN ID in a *switch-scoped VLAN* or with multiple (Port,VLAN) pairs in a *port-scoped VLAN*.
Distributed routing	All VTEPs in a virtual network perform inter-subnet routing and serve as L3 gateways in two possible modes: • Asymmetric routing: All VTEPs can perform routing, but routing is performed only on ingress VTEPs. Egress VTEPs perform bridging. • Symmetric routing: All VTEPs perform routing. Routing decisions are made on both ingress and egress VTEPs.
Virtual network	In OS10, each Layer 2 flooding domain in the overlay network is represented as a *virtual network*.
Virtual network identifier (VNID)	A 16-bit ID number that identifies a virtual network in OS10.
Virtual-network interface	A router interface that connects a virtual network bridge to a tenant VRF routing instance.
Access port	A port on a VTEP switch that connects to an end host and is part of the overlay network.
Network port	A port on a VTEP switch that connects to the underlay network.
Switch-scoped VLAN	A VLAN that is mapped to a virtual network ID (VNID) in OS10. All member ports of the VLAN are automatically added to the virtual network. • You can map only one VLAN ID to a virtual network. • Ideally suited for existing tenant VLANs that need to stretch over an IP fabric using VXLAN
Port-scoped VLAN	A Port,VLAN pair that is mapped to a virtual network ID (VNID) in OS10. You can assign an individual member interface to a virtual network either with an associated tagged VLAN or as an untagged member. Using a port-scoped VLAN, you can configure: • The same VLAN ID on different access interfaces to different virtual networks. • Different VLAN IDs on different access interfaces to the same virtual network.

VXLAN as NVO solution

Network virtualization overlay (NVO) is a solution designed to address the requirements of a multi-tenant data center, especially one with virtualized hosts. An NVO network is an overlay network used to extend L2 connectivity among VMs belonging to a tenant segment over an underlay IP network. Each tenant's payload is encapsulated in an IP packet at the originating VTEP and stripped of the encapsulation at the destination VTEP to access the payload. Each tenant segment is called a *virtual-network* and is uniquely identified in OS10 using a virtual network ID (VNID).

VXLAN is a type of encapsulation that is used as an NVO solution. VXLAN encapsulates a tenant's payload into IP UDP packets for transport across the IP underlay network. In OS10, each virtual network is assigned a 24-bit number called a *VXLAN network identifier* (VNI) that is carried in the VXLAN-encapsulated packet. The VNI uniquely identifies the tenant segment on all VTEPs. OS10 sets up ASIC tables to:

- Enable the creation of a Layer 2 bridge flooding domain across a Layer 3 network.
- Facilitate packet forwarding between local ports and tunneling packets from the local to a remote device.

Configure VXLAN

To extend a L2 tenant segment using VXLAN, follow these configuration steps on each VTEP switch:

1 Configure the source IP address used in encapsulated VXLAN packets.
2 Configure a virtual network and assign a VXLAN VNI.
3 Configure VLAN-tagged access ports.
4 Configure untagged access ports.
5 (Optional) Enable routing for hosts on different virtual networks.
6 Advertise the local VXLAN source IP address to remote VTEPs.
7 (Optional) Configure VLT.

Configure source IP address on VTEP

When you configure a switch as a VXLAN tunnel endpoint (VTEP), configure a loopback interface, whose IP address is used as the source IP address in encapsulated packet headers. Only a loopback interface assigned to a network virtualization edge (NVE) instance can be used as a source VXLAN interface.

- You cannot reconfigure the VXLAN source interface or the IP address assigned to the source interface if there is at least one VXLAN network ID (VNI) already assigned to a virtual-network ID (VNID) on the switch.
- The source loopback IP address must be reachable from a remote VTEP.
- An IPv6 address is not supported as the source VXLAN address.
- Do not assign the source loopback interface to a non-default VRF instance.
- Underlay reachability of remote tunnel endpoints is supported only in the default VRF.

1 Configure a loopback interface to serve as the source VXLAN tunnel endpoint in CONFIGURATION mode (0 to 255).
```
interface loopback number
```
2 Configure an IP address on the loopback interface in INTERFACE mode. The IP address allows the source VTEP to send VXLAN frames over the L3 transport network.
```
ip address ip-address/mask
```
3 Return to CONFIGURATION mode.
```
exit
```
4 Enter NVE mode from CONFIGURATION mode. NVE mode allows you to configure the VXLAN tunnel endpoint on the switch.
```
nve
```
5 Configure the loopback interface as the source tunnel endpoint for all virtual networks on the switch in NVE mode.
```
source-interface loopback number
```
6 Return to CONFIGURATION mode.
```
exit
```

Configure a VXLAN virtual network

To create a VXLAN, assign a VXLAN segment ID (VNI) to a virtual network ID (VNID) and configure a remote VTEP. A virtual network is identified by a unique 2-byte VNID. You cannot assign the same VXLAN VNI to more than one virtual network. VXLAN tunnel endpoints can either be manually configured in a static VXLAN or automatically discovered using BGP EVPN.

1 Create a virtual-network bridge domain in CONFIGURATION mode. Valid virtual-network ID (VNID) numbers are 1 to 65535.

```
virtual-network vn-id
```

2 Assign a VXLAN network identifier (VNI) to the virtual network in VIRTUAL-NETWORK mode (1 to 16,777,215). Configure the VNI for the same tenant segment on each VTEP switch.

```
vxlan-vni vni
```

3 (Optional) If you use BGP EVPN for VXLAN, this step is not required — To set up a static VXLAN, configure the source IP address of a remote VTEP in VXLAN-VNI mode. You can configure up to 1024 remote VTEP addresses for a VXLAN VNI.

```
remote-vtep ip-address
```

After you configure the remote VTEP, when the IP routing path to the remote VTEP IP address in the underlay IP network is known, the virtual network is enabled to send and receive VXLAN-encapsulated traffic from and to downstream servers and hosts. All broadcast, multicast, and unknown unicast (BUM) traffic received on access interfaces is replicated and sent to all configured remote VTEPs. Each packet contains the VXLAN VNI in its header.

By default, MAC learning from a remote VTEP is enabled and unknown unicast packets are flooded to all remote VTEPs. Re-enter the `remote-vtep ip-address` command to configure additional remote VTEPs.

4 Return to VIRTUAL-NETWORK mode.

```
exit
```

5 Return to CONFIGURATION mode.

```
exit
```

Configure VLAN-tagged access ports

Configure local access ports in the VXLAN overlay network using either a switch-scoped VLAN or port-scoped VLAN. Only one method is supported. You cannot assign tagged VLAN member interfaces to a virtual network using both switch-scoped and port-scoped VLANs.

- To use a switch-scoped VLAN to add VLAN-tagged member ports to a virtual network:
  a Assign a VLAN to the virtual network in VLAN Interface mode.

```
interface vlan vlan-id
virtual-network vn-id
```

  b Configure port interfaces as trunk members of the VLAN in Interface mode.

```
interface ethernet node/slot/port[:subport]
switchport mode trunk
switchport trunk allowed-vlan vlan-id
exit
```

  The local physical ports assigned to the VLAN transmit packets over the virtual network.

  NOTE: A switch-scoped VLAN assigned to a virtual network cannot have a configured IP address and cannot participate in L3 routing; for example:

```
OS10(config)# interface vlan 102
OS10(conf-if-vlan-5)# ip address 1.1.1.1/24
% Error: vlan102, IP address cannot be configured for VLAN attached to Virtual
Network.
```

- To use a port-scoped VLAN to add VLAN-tagged member ports to a virtual network:

a Configure interfaces as trunk members in Interface mode.

```
interface ethernet node/slot/port[:subport]
switchport mode trunk
exit
```

b Assign a trunk member interface as a Port,VLAN ID pair to the virtual network in VIRTUAL-NETWORK mode. All traffic sent and received for the virtual network on the interface carries the VLAN tag. Multiple tenants connected to different switch interfaces can have the same `vlan-tag` VLAN ID.

```
virtual-network vn-id
member-interface ethernet node/slot/port[:subport] vlan-tag vlan-id
```

The Port,VLAN pair starts to transmit packets over the virtual network.

c Repeat Steps a) and b) to assign additional member Port,VLAN pairs to the virtual network.

- You cannot assign the same Port,VLAN member interface pair to more than one virtual network.
- You can assign the same `vlan-tag` VLAN ID with different member interfaces to different virtual networks.
- You can assign a member interface with different `vlan-tag` VLAN IDs to different virtual networks.

The VLAN ID tag is removed from packets transmitted in a VXLAN tunnel. Each packet is encapsulated with the VXLAN VNI in the packet header before it is sent from the egress source interface for the tunnel. At the remote VTEP, the VXLAN VNI is removed and the packet is transmitted on the virtual-network bridge domain. The VLAN ID is regenerated using the VLAN ID associated with the virtual-network egress interface on the VTEP and included in the packet header.

Configure untagged access ports

Add untagged access ports to the VXLAN overlay network using either a switch-scoped VLAN or port-scoped VLAN. Only one method is supported.

- To use a switch-scoped VLAN to add untagged member ports to a virtual network:

  a Assign a VLAN to a virtual network in VLAN Interface mode.

  ```
  interface vlan vlan-id
  virtual-network vn-id
  exit
  ```

  b Configure port interfaces as access members of the VLAN in Interface mode.

  ```
  interface ethernet node/slot/port[:subport]
  switchport access vlan vlan-id
  exit
  ```

  Packets received on the untagged ports are transmitted over the virtual network.

- To use a port-scoped VLAN to add untagged member ports to a virtual network:

  a Create a reserved VLAN ID to assign untagged traffic on member interfaces to a virtual network in CONFIGURATION mode. The VLAN ID is used internally for all untagged member interfaces on the switch that belong to virtual networks.

  ```
  virtual-network untagged-vlan untagged-vlan-id
  ```

  b Configure port interfaces as trunk members and remove the access VLAN in Interface mode.

  ```
  interface ethernet node/slot/port[:subport]
  switchport mode trunk
  no switchport access vlan
  exit
  ```

  c Assign the trunk interfaces as untagged members of the virtual network in VIRTUAL-NETWORK mode. You cannot use the reserved VLAN ID for a legacy VLAN or for tagged traffic on member interfaces of virtual networks.

  ```
  virtual-network vn-id
  member-interface ethernet node/slot/port[:subport] untagged
  exit
  ```

If at least one untagged member interface is assigned to a virtual network, you cannot delete the reserved untagged VLAN ID. If you reconfigure the reserved untagged VLAN ID, you must either reconfigure all untagged member interfaces in the virtual networks to use the new ID or reload the switch.

Enable overlay routing between virtual networks

The previous sections described how a VTEP switches traffic between hosts within the same Layer 2 tenant segment (virtual network) and transports traffic over an IP underlay fabric. This section describes how a VTEP enables hosts in different Layer 2 segments belonging to the same tenant VRF to communicate with each other.

NOTE: On the S4248-ON, IPv6 overlay routing between virtual networks is not supported with static VXLAN. IPv6 overlay routing is, however, supported with BGP EVPN.

Each tenant is assigned a VRF and each virtual-network interface is assigned an IP subnet in the tenant VRF. The VTEP acts as the L3 gateway that routes traffic from one tenant subnet to another in the overlay before encapsulating it in the VXLAN header and transporting it over the IP underlay fabric.

To enable the routing of host traffic between virtual networks, configure an interface for each virtual network and associate it to a tenant VRF. Assign a unique IP address in the IP subnet range associated with the virtual network to each virtual-network interface on each VTEP.

To enable the most efficient traffic forwarding on a VTEP, OS10 supports distributed gateway routing. A distributed gateway means that multiple VTEPs can act as the gateway router for a tenant subnet. The VTEP that is located nearest to a host acts as its gateway router. To support seamless migration of hosts and virtual machines on different VTEPs, configure a common virtual IP address (anycast IP address) on all VTEPs for each virtual network. Use this anycast IP address as the gateway IP address on virtual machines.

To support multiple tenants when each tenant has its own set of Layer 2 segments, configure a different IP VRF for each tenant. All tenants share the same VXLAN underlay IP fabric in the default VRF.

1 As a best practice, create a non-default VRF instance for overlay routing in Configuration mode. For multi-tenancy, create a VRF instance for each tenant.

```
ip vrf tenant-vrf-name
exit
```

2 Configure the anycast gateway MAC address used by all VTEPs in all VXLAN virtual networks in Configuration mode.

When a virtual machine sends an ARP request for the anycast gateway IP address in a VXLAN virtual network, the nearest VTEP responds with the configured anycast MAC address. Configure the same MAC address on all VTEPs so that the anycast gateway MAC address remains the same if a virtual machine migrates to a different VTEP. Because the configured MAC address is automatically used for all VXLAN virtual networks, configure it in global Configuration mode.

```
ip virtual-router mac-address mac-address
```

3 Configure a virtual-network interface, assign it to the tenant VRF, and configure an IP address.

The interface IP address must be unique on each VTEP, including VTEPs in VLT pairs. You can configure an IPv6 address on the virtual-network interface. Different virtual-network interfaces configured on the same VTEP must have virtual-network IP addresses in different subnets. If you do not assign the virtual-network interface to a tenant VRF, it is assigned to the default VRF.

```
interface virtual-network vn-id
ip vrf forwarding tenant-vrf-name
ip address ip-address/mask
no shutdown
exit
```

4 Configure an anycast gateway IPv4 or IPv6 address for each virtual network in INTERFACE-VIRTUAL-NETWORK mode. This anycast IP address must be in the same subnet as the IP address of the virtual-network interface in Step 3.

Configure the same IPv4 or IPv6 address as the anycast IP address on all VTEPs in a virtual network. The anycast gateway IP address is used as the default gateway IP address by all hosts in the subnet that is connected to the virtual-network interface configured in Step 3. Configure the anycast gateway IP address on all downstream VMs. Using the same anycast gateway IP address allows host VMs to move from one VTEP to another VTEP in a VXLAN. Using an anycast gateway is the recommended best practice in both VLT and non-VLT VXLAN configurations.

```
interface virtual-network vn-id
ip virtual-router address ip-address
```

Configuration notes for virtual-network routing:

- VXLAN overlay routing includes routing tenant traffic on the ingress VTEP and bridging the traffic on the egress VTEP. The ingress VTEP learns ARP entries and associates all destination IP addresses of tenant VMs with the corresponding VM MAC addresses in the overlay. On the ingress VTEP, configure a virtual network for each destination IP subnet even if there are no locally attached hosts for an IP subnet.
- Routing protocols, such as OSPF and BGP, are not supported on the virtual-network interface in the overlay network. However, static routes that point to a virtual-network interface or to a next-hop IP address that belongs to a virtual-network subnet are supported.
- When you add a static route in the overlay, any next-hop IP address that belongs to a virtual-network subnet must be the only next-hop for that route and cannot be one of multiple ECMP next-hops. For example, if you enter the following configuration commands one after the other, where the 10.250.0.0/16 is a virtual-network subnet, only the first next-hop is active on the switch.

  ```
  OS10(config)# ip route 0.0.0.0/0  10.250.0.101
  OS10(config)# ip route 0.0.0.0/0  10.250.0.102
  ```

  If the next-hop is a pair of dual-homed VTEPs in a VLT domain, a workaround is to configure the same anycast gateway IP address on both VTEPs and use this address as the next-hop IP address.
- VLT peer routing is not supported in a virtual network. A packet destined to the virtual-network peer MAC address is L2 switched instead of IP routed. To achieve active-active peer routing in a virtual network, configure the same virtual anycast gateway IP and MAC addresses on both VTEP VLT peers and use the anycast IP as default gateway on the VMs.
- VRRP is not supported on a virtual-network interface. Configure the virtual anycast gateway IP address to share a single gateway IP address on both VTEP VLT peers and use the anycast IP as default gateway on the VMs.
- IGMP and PIM are not supported on a virtual-network interface.
- IP routing of incoming VXLAN encapsulated traffic in the overlay after VXLAN termination is not supported.

The following tables show an example of how to use anycast gateway IP and MAC addresses in a data center with three virtual networks and multiple VTEPs:

- Globally configure an anycast MAC address for all VTEPs in all virtual networks. For example, if you use three VTEP switches in three virtual networks:

Table 17. MAC address for all VTEPs

Virtual network	VTEP	Anycast gateway MAC address
VNID 11	VTEP 1	00.11.22.33.44.55
	VTEP 2	00.11.22.33.44.55
	VTEP 3	00.11.22.33.44.55
VNID 12	VTEP 1	00.11.22.33.44.55
	VTEP 2	00.11.22.33.44.55
	VTEP 3	00.11.22.33.44.55
VNID 13	VTEP 1	00.11.22.33.44.55
	VTEP 2	00.11.22.33.44.55
	VTEP 3	00.11.22.33.44.55

- Configure a unique IP address on the virtual-network interface on each VTEP across all virtual networks. Configure the same anycast gateway IP address on all VTEPs in a virtual-network subnet. For example:

Table 18. IP address on the virtual-network interface on each VTEP

Virtual network	VTEP	Virtual-network IP address	Anycast gateway IP address
VNID 11	VTEP 1	10.10.1.201	10.10.1.254
	VTEP 2	10.10.1.202	10.10.1.254

Virtual network	VTEP	Virtual-network IP address	Anycast gateway IP address
	VTEP 3	10.10.1.203	10.10.1.254
VNID 12	VTEP 1	10.20.1.201	10.20.1.254
	VTEP 2	10.20.1.202	10.20.1.254
	VTEP 3	10.20.1.203	10.20.1.254
VNID 13	VTEP 1	10.30.1.201	10.30.1.254
	VTEP 2	10.30.1.202	10.30.1.254
	VTEP 3	10.30.1.203	10.30.1.254

Advertise VXLAN source IP address

1 Advertise the IP address of the local source tunnel interface to all VTEPs in the underlay IP network using the existing routing infrastructure. This example uses OSPF to advertise the VXLAN source IP address on Ethernet1/1/3, which is the underlay network-facing interface:

```
OS10(config)# router ospf 100
OS10(config-ospf)# router-id 110.111.170.195
OS10(config-ospf)# exit
OS10(config)# interface ethernet1/1/3
OS10(config-if-eth1/1/3)# ip ospf 100 area 0.0.0.0
OS10(config-if-eth1/1/3)# exit
OS10(config)# interface loopback 1
OS10(config-if-lo-1)# ip ospf 100 area 0.0.0.0
```

Each VTEP switch in the underlay IP network learns the IP address of the VXLAN source interface. If a remote VTEP switch is not reachable, its status displays as `DOWN` in `show nve remote-vtep` output.

2 Configure the MTU value on L3 underlay network-facing interfaces in Interface mode (1312 to 9216) to be at least 50 bytes higher than the MTU on the server-facing links to allow for VXLAN encapsulation .

```
mtu value
```

3 Return to CONFIGURATION mode.

```
exit
```

Configure VLT

(Optional) To use VXLAN in a VLT domain, configure the VLT domain — including the VLT Interconnect (VLTi) interfaces, backup heartbeat, and VLT MAC address — as described in Virtual link trunking.

Required VLT VXLAN configuration:

- The IP address of the VTEP source loopback interface must be same on the VLT peers.
- If you use a port-scoped VLAN to assign tagged access interfaces to a virtual network, you must configure a unique VLAN ID for the VLT Interconnect (VLTi) link to identify traffic belonging to each virtual network. Configure a VLAN to transmit VXLAN traffic over the VLTi link in VIRTUAL-NETWORK mode. All traffic sent and received from a virtual network on the VLTi carries the VLTi VLAN ID tag.

  Configure the same VLTi VLAN ID on both VLT peers. You cannot use the ID of an existing VLAN on a VLT peer or the reserved untagged VLAN ID. You can use the VLTi VLAN ID to assign tagged or untagged access interfaces to a virtual network.

```
virtual-network vn-id
vlti-vlan vlan-id
```

- Although a VXLAN virtual network has no access port members that connect to downstream servers, you must still configure a switch-scoped VLAN or VLTi VLAN. The presence of this VLAN ensures that the VLTi link is added as a member of the virtual network so that mis-hashed ARP packets received from the VXLAN tunnel reach the intended VLT node.

Best practices:

- If a VLT peer loses connectivity to the underlay L3 network, it continues to transmit routing traffic to the network through the VLTi link on a dedicated L3 VLAN to the other VLT peer. It is best practice to configure a L3 VLAN between VLT peers in the underlay network and enable routing on the VLAN; for example:

```
OS10(config)# interface vlan4000
OS10(config-if-vl-4000)# no shutdown
OS10(config-if-vl-4000)# ip address 41.1.1.1/24
OS10(config-if-vl-4000)# ip ospf 1 area 0.0.0.0
```

- To reduce traffic loss when a VLT peer boots up and joins an existing VLT domain, or when the VLTi links fails and the VLT peer is still up as detected by the VLT heartbeat, create an uplink state group. Configure all access VLT port channels on the peer as upstream links. Configure all network-facing links as downstream link. For example:

```
OS10(config)# uplink-state-group 1
OS10(conf-uplink-state-group-1)# enable
OS10(conf-uplink-state-group-1)# downstream ethernet1/1/1-1/1/2
OS10(conf-uplink-state-group-1)# upstream port-channel 10
```

L3 VXLAN route scaling

The S41xx-ON and S52xx-ON series, S4048T-ON, and S6010-ON switches support native VxLAN routing — routing in and out of tunnels (RIOT), RIOT requires dedicated hardware resources to be reserved for overlay routing. These dedicated resources cannot be used for underlay routing.

Each overlay ARP entry requires a routing next-hop in the hardware to bind a destination tenant VM IP address to the corresponding tenant VM MAC address and VNI. Each virtual-network interface that is assigned to an IP subnet requires a routing interface in the hardware.

OS10 supports preset profiles to re-allocate the number of resources reserved for overlay ARP entries. The number of entries reserved for each preset mode differs according to OS10 switch.

Table 19. Routing next-hops reserved on OS10 switches

OS10 Switch	Overlay next-hop entries	Underlay next-hop entries	Overlay L3 RIF entries	Underlay L3 RIF entries
S41xx-ON series:	4096	28672	2048	10240
default-overlay-routing	0	32768	0	12288
disable-overlay-routing	16384	16384	6144	6144
balanced-overlay-routing	24576	8192	10240	2048
scaled-overlay-routing				
S4048T-ON and S6010-ON:	8192	4096	2048	2048
default-overlay-routing	0	49152	49152	0
disable-overlay-routing	24576	24576	24576	6144
balanced-overlay-routing	40960	8192	8192	10240
scaled-overlay-routing				
S52xx-ON series:	8192	57344	2048	14336

OS10 Switch	Overlay next-hop entries	Underlay next-hop entries	Overlay L3 RIF entries	Underlay L3 RIF entries
default-overlay-routing	0	65536	0	16384
disable-overlay-routing	32768	32768	8192	8192
balanced-overlay-routing	53248	12288	12288	4096
scaled-overlay-routing				
S4248-ON: default-overlay-routing	20480	110592	4096	28672

NOTE: The S4248-ON supports only one default profile to reserve resources for overlay ARP entries.

After you configure an overlay routing profile, save the configuration and reload the switch to activate the profile.

Configure an overlay routing profile

- Enable an overlay routing profile in Configuration mode or disable the configured profile and return to the default.

```
OS10# hardware overlay-routing-profile {disable-overlay-routing | balanced-overlay-routing |
scaled-overlay-routing}
```

Display overlay routing profiles

- View the hardware resources available for overlay routing in different profiles; for example, S52xx-ON series:

```
OS10# show hardware overlay-routing-profile mode all
Mode                      Overlay Next-hop  Underlay Next-hop  Overlay L3 RIF  Underlay L3 RIF
                           Entries           Entries            Entries         Entries
default-overlay-routing   8192              57344              2048            14336
disable-overlay-routing   0                 65536              0               16384
balanced-overlay-routing  32768             32768              8192            8192
scaled-overlay-routing    53248             12288              12288           4096
```

- View the currently configured overlay routing profile; for example, S52xx-ON series:

```
show hardware overlay-routing-profile mode
                                  Overlay    Underlay   Overlay   Underlay
Setting     Mode                  Next-hop   Next-hop   L3 RIF    L3 RIF
                                  Entries    Entries    Entries   Entries
Current    default-overlay-routing   8192     57344      2048      14336
Next-boot  default-overlay-routing   8192     57344      2048      14336
```

DHCP relay on VTEPs

DHCP clients on hosts in the overlay can communicate with a DHCP server using DHCP relay on the VTEP switch. In OS10, DHCP relay is supported on VTEPs only if the DHCP server is located in the underlay network. To work seamlessly, VTEP DHCP relay transmits the virtual-network IP address of the relay interface to the DHCP server.

By default, DHCP uses the `giaddr` packet field to carry these addresses to the server. In a VxLAN which has overlay and underlay subnets in the same default VRF, DHCP relay on VTEPs operates without user intervention. However, in a VXLAN in which the underlay and overlay are in different VRFs, the default DHCP method is not successful. The IP tenant subnet is in the overlay address space. The IP address on which the VTEP is reachable is in the underlay address space. You must configure an additional DHCP sub-option (5 or 151) in DHCP relay agent option 82 to transmit the IP subnet of the client separately from the IP address on which the VTEP is reachable.

Because OS10 does not support the required sub-options in DHCP relay agent option 82, the `giaddr` packet field must contain the virtual-network IP address of the relay interface, and this IP address must be reachable from the DHCP server in the underlay. Each VTEP that acts as a DHCP relay must have its virtual-network IP address installed (via a route leaking mechanism) as a route to the underlay and advertised to all underlay routers, including the spine switches.

Similarly, the DHCP server in the underlay VRF must be reachable from the client tenant VRF in the overlay. Configure a static route for DHCP server subnet in the underlay default VRF, and leak the static route to the client tenant VRF in the overlay. This configuration sets up

a bi-directional communication between the client and DHCP server across the virtual networks. The route-leaking configuration is not required if the VxLAN overlay subnet and underlay subnet are in same default VRF.

Configure DHCP relay on VTEPs

1 Configure the IP address of the virtual-network relay interface in the non-default tenant VRF as a static route in the default VRF; for example:

```
OS10(config)# ip route 10.10.0.2/32 interface virtual-network 10
```

2 Configure a static IP route to the DHCP server interface in the tenant VRF; for example:

```
OS10(config)# ip route vrf tenant01 40.1.1.0/24 interface vlan40
```

3 Configure DHCP relay on the virtual-network interface of the tenant VRF; for example:

```
OS10(config)# interface virtual-network 10
OS10(conf-if-vn-10)# ip helper-address 40.1.1.1 vrf tenant01
```

View VXLAN configuration

Use `show` commands to verify the VXLAN configuration and monitor VXLAN operation.

View VXLAN virtual network

```
OS10# show virtual-network
Codes: DP - MAC-learn Dataplane, CP - MAC-learn Controlplane, UUD - Unknown-Unicast-Drop
Un-tagged VLAN: 888
Virtual Network: 60000
   VLTi-VLAN: 2500
   Members:
      VLAN 1000: port-channel1, ethernet1/1/9, ethernet1/1/10
      VLAN 2500: port-channel1000
   VxLAN Virtual Network Identifier: 16775000
      Source Interface: loopback100(222.222.222.222)
      Remote-VTEPs (flood-list): 55.55.55.55(DP),77.1.1.1(DP)
```

View VXLAN virtual-network port

```
OS10# show virtual-network interface ethernet 1/1/1
Interface          Vlan      Virtual-network
ethernet1/1/1      100         1000
ethernet1/1/1      200         2000
ethernet1/1/1      300         3000
```

View VXLAN virtual-network VLAN

```
OS10# show virtual-network vlan 100
Vlan  Virtual-network   Interface
100    1000             ethernet1/1/1,ethernet1/1/2
100    5000             ethernet1/1/2
```

View VXLAN virtual-network VLANs

```
OS10# show vlan
Codes: * - Default VLAN, M - Management VLAN, R - Remote Port Mirroring VLANs,
@ - Attached to Virtual Network
Q: A - Access (Untagged), T - Tagged

 NUM   Status Description Q Ports
* 1      up               A Eth1/1/1-1/1/48
@ 100    up               T Eth1/1/2,Eth1/1/3
                          A Eth1/1/1
@ 101    up               T port-channel5
  200    up               T Eth1/1/11-1/1/15
```

View VXLAN virtual-network statistics

```
OS10# show virtual-network counters
Virtual-Network      Input (Packets/Bytes)       Output (Packets/Bytes)
1000                 857/8570                    257/23709
2000                 457/3570                    277/13709
```

```
OS10# show virtual-network counters interface 1/1/3 vlan 100
Virtual-Network      Input (Packets/Bytes)       Output (Packets/Bytes)
1000                 857/8570                    257/23709
2000                 457/3570                    277/13709
```

NOTE: Using flex counters, OS10 may display additional packets in the Output field number, but the additional packets are not transmitted. For an accurate count, use the Output Bytes number.

View VXLAN remote VTEPs

```
OS10# show nve remote-vtep summary
Remote-VTEP      State
----------------------
2.2.2.2          up
```

```
OS10# show nve remote-vtep
Codes: DP - MAC-learn Dataplane, CP - MAC-learn Controlplane, UUD - Unknown-Unicast-Drop
IP Address: 2.2.2.2,  State: up, Encap: VxLAN
   VNI list: 10000(DP), 200(DP), 300(DP)
```

View VXLAN statistics on remote VTEPs

```
OS10# show nve remote-vtep counters
Remote-VTEP          Input (Packets/Bytes)       Output (Packets/Bytes)
-------------------------------------------------------------------------
10.10.10.10          857/8570                    257/23709
20.20.20.20          457/3570                    277/13709
```

View VXLAN virtual network by VN-ID

```
OS10# show nve vxlan-vni
VNI     Virtual-Network    Source-IP      Remote-VTEPs
-----------------------------------------------------
101     101                44.44.44.44    11.11.11.11,22.22.22.22,33.33.33.33
102     102                44.44.44.44    11.11.11.11,22.22.22.22,33.33.33.33
103     103                44.44.44.44    11.11.11.11,22.22.22.22,33.33.33.33
104     104                44.44.44.44    11.11.11.11,22.22.22.22,33.33.33.33
```

View VXLAN routing between virtual networks

The `show ip arp vrf` and `show ipv6 neighbors vrf` outputs display information about IPv4 and IPv6 neighbors learned in a non-default VRF on the switch. The `show ip route vrf` command displays the IPv4 and IPv6 routes learned.

```
OS10# show ip arp vrf tenant1
Address    Hardware address   Interface         Egress Interface
-------------------------------------------------------------------
111.0.0.2  00:c5:15:02:12:f1  virtual-network20 ethernet1/1/5
111.0.0.3  00:c5:15:02:12:a2  virtual-network20 port-channel5
111.0.0.4  00:12:98:1f:34:11  virtual-network20 VXLAN(20.0.0.1)
121.0.0.3  00:12:28:1f:34:15  virtual-network20 port-channel5
121.0.0.4  00:f2:34:ac:34:09  virtual-network20 VXLAN(20.0.0.1)

OS10# show ipv6 neighbors vrf tenant1
IPv6 Address    Hardware Address   State  Interface         Egress Interface
-----------------------------------------------------------------------------
200::2          00:12:28:1f:34:15  STALE  virtual-network40 port-channel5
200::f          00:f2:34:ac:34:09  REACH  virtual-network40 VXLAN(20.0.0.1)
```

```
OS10# show ip route vrf vrf_1
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
```

```
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, * - candidate default,
       + - summary route, > - non-active route
Gateway of last resort is not set
  Destination       Gateway         Dist/Metric              Last Change
-------------------------------------------------------------------------
  C  100.1.0.0/16   via 100.1.1.4   virtual-network60000  0/0  00:36:24
  C  100.33.0.0/16  via 100.33.1.4  virtual-network60032  0/0  00:36:23
  C  100.65.0.0/16  via 100.65.1.4  virtual-network60064  0/0  00:36:22
  C  100.97.0.0/16  via 100.97.1.4  virtual-network60096  0/0  00:36:21

OS10# show ipv6 route vrf vrf_1
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, * - candidate default,
       + - summary route, > - non-active route
Gateway of last resort is not set
   Destination         Gateway               Dist/Metric              Last Change
-------------------------------------------------------------------------------
 C 1000:100:10:1::/64  via 1000:100:10:1::4  virtual-network60000 0/0 00:37:08
 C 1000:100:10:21::/64 via 1000:100:10:21::4 virtual-network60032 0/0 00:37:07
 C 1000:100:10:41::/64 via 1000:100:10:41::4 virtual-network60064 0/0 00:37:06
 C 1000:100:10:61::/64 via 1000:100:10:61::4 virtual-network60096 0/0 00:37:05
```

VXLAN MAC addresses

Use the `show mac address-table virtual-network` or `show mac address-table extended` commands to display the MAC addresses learned on a VXLAN virtual network or learned on both VXLAN virtual networks and legacy VLANs.

Use the `clear mac address-table dynamic virtual-network` and `clear mac address-table dynamic nve remote-vtep` commands to delete address entries from the MAC address virtual-network table.

NOTE: The existing `show mac address-table` and `clear mac-address table` commands do not display and clear MAC addresses in a virtual-network bridge domain even when access ports in a switch-scoped VLAN are assigned to a VXLAN virtual network.

Display VXLAN MAC addresses

Table 20. Display VXLAN MAC addresses

Command	Description
`show mac address-table virtual-network [`*`vn-id`*` | local | remote | static | dynamic | address `*`mac-address`*` | interface {ethernet `*`node/slot/port:subport`*` | port-channel `*`number`*`}]`	Displays all MAC addresses learned on all or a specified virtual network. *`vn-id`*: Displays only information about the specified virtual network. `local`: Displays only locally-learned MAC addresses. `remote`: Displays only remote MAC addresses. `static`: Displays only static MAC addresses. `dynamic`: Displays only dynamic MAC addresses. `address` *`mac-address`*: Displays only information about the specified MAC address.

Command	Description
	`interface ethernet` *node/slot/port:subport*: Displays only MAC addresses learned on the specified interface. `interface port-channel` *number*: Displays only MAC addresses learned on the specified port channel.
`show mac address-table extended [address` *mac-address* `| interface {ethernet` *node/slot/port:subport* `| port-channel` *number*`} | static | dynamic]`	Displays MAC addresses learned on all VLANs and VXLANs (default). `address` *mac-address*: Displays only information about the specified MAC address. `interface ethernet` *node/slot/port:subport*: Displays only MAC addresses learned on the specified interface. `interface port-channel` *number*: Displays only MAC addresses learned on the specified port channel. `static`: Displays only static MAC addresses. `dynamic`: Displays only dynamic MAC addresses.
`show mac address-table nve {vxlan-vni` *vn-id* `| remote-vtep` *ip-address*`}`	`vxlan-vni` *vn-id*: Displays MAC addresses learned on NVE from the specified VXLAN virtual-network ID. `remote-vtep` *ip-address*: Displays MAC addresses learned on NVE from the specified remote VTEP.
`show mac address-table count virtual-network [dynamic | local | remote | static | interface {ethernet` *node/slot/port:subport* `| port-channel` *number*`} |` *vn-id*`]`	Displays the number of MAC addresses learned on all virtual networks (default). `dynamic`: Displays the number of dynamic MAC addresses learned on all or a specified virtual network. `local`: Displays the number of locally-learned MAC addresses. `remote`: Displays the number of remote MAC addresses learned on all or a specified virtual network. `static`: Displays the number of static MAC addresses learned on all or a specified virtual network. `interface ethernet` *node/slot/port:subport*: Displays the number of MAC addresses learned on the specified interface. `interface port-channel` *number*: Displays the number of MAC addresses learned on the specified port channel. *vn-id*: Displays the number of MAC addresses learned on the specified virtual network.
`show mac address-table count nve {remote-vtep` *ip-address* `| vxlan-vni` *vn-id*`}`	Displays the number of MAC addresses learned for a virtual network or from a remote VTEP. `remote-vtep` *ip-address*: Displays the number of MAC addresses learned on the specified remote VTEP. `vxlan-vni` *vn-id*: Displays the number of MAC addresses learned on the specified VXLAN virtual network.

Command	Description
`show mac address-table count extended [interface ethernet node/slot/port:subport | port-channel number]}`	Displays the number of MAC addresses learned on all VLANs and VXLAN virtual networks. `interface ethernet node/slot/port:subport`: Displays the number of MAC addresses learned from VLANs and VXLANs on the specified interface. `port-channel number`: Displays the number of MAC addresses learned from VLANs and VXLANs on the specified port channel.

Clear VXLAN MAC addresses

Table 21. Clear VXLAN MAC addresses

Command	Description
`clear mac address-table dynamic virtual-network [interface {ethernet node/slot/port:subport | port-channel number} | local | vn-id [address mac-address | local]]`	Clears all MAC addresses learned on all VXLAN virtual networks. `interface ethernet node/slot/port:subport`: Clears only MAC addresses learned on the specified interface. `interface port-channel number`: Clears only MAC addresses learned on the specified port channel. `local`: Clears only locally-learned MAC addresses. `vn-id`: Clears only the MAC addresses learned on the specified virtual network. `vn-id address mac-address`: Clears only the MAC address learned on the specified virtual network.
`clear mac address-table dynamic nve remote-vtep ip-address`	Clears all MAC addresses learned from the specified remote VTEP.

VXLAN commands

hardware overlay-routing-profile

Configures the number of reserved ARP table entries for VXLAN overlay routing.

Syntax `hardware overlay-routing-profile {balanced-overlay-routing | scaled-overlay-routing | disable-overlay-routing}`

Parameters

balanced-overlay-routing Reserve routing entries for balanced VXLAN tenant routing:

- S4048T-ON and S6010-ON: 24576 entries
- S41xx-ON series: 16384 entries
- S52xx-ON series switches: 32768 entries

scaled-overlay-routing Reserve routing entries for scaled VXLAN tenant routing:

- S4048T-ON and S6010-ON: 36864 entries

- S41xx-ON series: 24576 entries
- S52xx-ON series switches: 53248 entries

disable-overlay-routing	Allocate 0 next-hop entries for overlay routing and all next-hop entries for underlay routing.

Default S4048T-ON and S6010-ON switches reserve 8192 ARP table entries.

S41xx-ON series switches reserve 4096 ARP table entries.

S52xx-ON series switches reserve 8192 ARP table entries.

Command mode CONFIGURATION

Usage information The number of reserved table entries in a profile varies according to the OS10 switch. To view the available overlay routing profiles for a switch, use the `show hardware overlay-routing-profile mode all` command. After you configure a profile, reload the switch to activate the profile. The `no` form of the command disables the configured profile and restores the default number of reserved ARP table entries.

Example

```
OS10(config)# hardware overlay-routing-profile balanced-overlay-routing
OS10(config)# exit
OS10# write memory
OS10# reload
```

Supported releases 10.4.3.0 or later

interface virtual-network

Configures a virtual-network router interface.

Syntax `interface virtual-network vn-id`

Parameters

virtual-network *vn-id*	Enter a virtual-network ID, from 1 to 65535.

Default Not configured

Command mode CONFIGURATION

Usage information Configure a virtual-network router interface to enable hosts connected to a virtual network to route traffic to hosts on another virtual network in the same VRF. The virtual-network IP address must be unique on each VTEP, including VTEPs in VLT pairs.

Example

```
OS10(config)# interface virtual-network 10000
OS10(config-if-vn-10000)# ip vrf forwarding tenant1
OS10(config-if-vn-10000)# ip address 10.1.0.1/16
OS10(config-if-vn-10000)# no shutdown
```

Supported releases 10.4.3.0 or later

ip virtual-router address

Configures an anycast gateway IP address for a VXLAN virtual network.

Syntax `ip virtual-router address ip-address`

Parameters

address *ip-address*	Enter the IP address of the anycast Layer 3 gateway.

Default Not configured

Command mode INTERFACE-VIRTUAL-NETWORK

Usage information Configure the same anycast gateway IP address on all VTEPs in a VXLAN virtual network. The anycast gateway IP address is used as the default gateway IP address if host VMs move from one VTEP to another in a VXLAN. The anycast gateway IP address must be in the same subnet as the IP address of the virtual-network router interface.

Example

```
OS10(config)# interface virtual-network 10000
OS10(config-if-vn-10000)# ip virtual-router address 10.1.0.100
```

Supported releases 10.4.3.0 or later

ip virtual-router mac-address

Configures the MAC address of an anycast Layer 3 gateway for VXLAN routing.

Syntax `ip virtual-router mac-address mac-address`

Parameters

mac-address *mac-address*	Enter the MAC address of the anycast Layer 3 gateway.

Default Not configured

Command mode CONFIGURATION

Usage information Configure the same MAC address on all VTEPs so that the anycast gateway MAC address remains the same if a virtual machine migrates to a different VTEP. Because the configured MAC address is automatically used for all VXLAN virtual networks, you configure it in global Configuration mode.

Example

```
OS10(config)# ip virtual-router mac-address 00:01:01:01:01:01
```

Supported releases 10.4.3.0 or later

member-interface

Assign untagged or tagged VLAN traffic on a member interface to a virtual network.

Syntax

```
member-interface {ethernet node/slot/port[:subport] | port-channel number}
{vlan-tag vlan-id | untagged}
```

Parameters	**ethernet *node/slot/port[:subport]***	Assign the specified interface to a virtual network.
	port-channel *number*	Assign the specified port channel to a virtual network.
	untagged	Assign untagged traffic on an interface or port channel to a virtual network.
	vlan-tag *vlan-id*	Assign tagged traffic on the specified VLAN to a virtual network.

Default Not configured

Command mode VIRTUAL-NETWORK

Usage information Use the `member-interface` command to assign traffic on the same VLAN or interface to different virtual networks. The `no` version of this command removes the configured value.

Example

```
OS10(config)# virtual-network 10000
OS10(config-vn)# member-interface port-channel 10 vlan-tag 200
OS10(config-vn)# member-interface port-channel 20 untagged
```

Supported releases 10.4.2.0 or later

nve

Enter network virtualization edge (NVE) configuration mode to configure source VXLAN tunnel endpoint.

Syntax `nve`

Parameters None

Default None

Command mode CONFIGURATION

Usage information In NVE mode, you can configure the source tunnel endpoint for all virtual networks on the switch.

Example

```
OS10# nve
OS10(config-nve)#
```

Supported releases 10.4.2.0 or later

remote-vtep

Configures the IP address of a remote tunnel endpoint in a VXLAN network.

Syntax `remote-vtep` *`ip-address`*

Parameters *`ip-address`* — Enter the IP address of a remote virtual tunnel endpoint (VTEP).

Default Not configured

Command mode VIRTUAL-NETWORK VXLAN-VNI

Usage information After you configure the remote VTEP, the VXLAN virtual network is enabled to start sending server traffic. You can configure multiple remote VTEPs. All broadcast, multicast, and unknown unicast (BUM) traffic received on an access interface is replicated on remote VTEPs. The `no` version of this command removes the configured value.

Example

```
OS10(config-vn-vxlan-vni)# remote-vtep 20.20.20.1
OS10(config-vn-vxlan-vni-remote-vtep)# exit
OS10(config-vn-vxlan-vni)# remote-vtep 30.20.20.1
```

Supported releases 10.4.2.0 or later

show hardware overlay-routing-profile mode

Displays the number of hardware resources available for overlay routing in different profiles.

Syntax `show hardware overlay-routing-profile mode [all]`

Parameters

all	View the number of tenant entries available in each hardware partition for overlay routing profiles.

Default Not configured

Command mode EXEC

Usage information On S41xx-ON, and S52xx-ON series, S4048T-ON, S4248-ON, and S6010-ON switches, L3 VXLAN overlay routing requires hardware resources to be reserved. The number of reserved table entries in a profile varies according to the OS10 switch.

Example (S52xx-ON series)

```
OS10# show hardware overlay-routing-profile mode all
                          Overlay    Underlay   Overlay   Underlay
Mode                      Next-hop   Next-hop   L3 RIF    L3 RIF
                          Entries    Entries    Entries   Entries
default-overlay-routing   8192       57344      2048      14336
disable-overlay-routing   0          65536      0         16384
balanced-overlay-routing  32768      32768      8192      8192
scaled-overlay-routing    53248      12288      12288     4096
```

```
show hardware overlay-routing-profile mode
                                    Overlay   Underlay  Overlay  Underlay
Setting     Mode                    Next-hop  Next-hop  L3 RIF   L3 RIF
                                    Entries   Entries   Entries  Entries
Current   default-overlay-routing   8192      57344     2048     14336
Next-boot default-overlay-routing   8192      57344     2048     14336
```

Supported releases 10.4.3.0 or later

show interface virtual-network

Displays the configuration of virtual-network router interfaces and packet statistics.

Syntax `show interface virtual-network [vn-id]`

Parameters

vn-id	Enter a virtual-network ID (1 to 65535).

Default Not configured

Command mode EXEC

Usage information Use the `show interface virtual-network` command to display the virtual-network IP address used for routing traffic in a virtual network. Traffic counters are also displayed.

Example

```
show interface virtual-network 102
Virtual-network 102 is up, line protocol is up
Address is 14:18:77:25:6f:84, Current address is 14:18:77:25:6f:84
Interface index is 66
Internet address is 12.12.12.2/24
Mode of IPv4 Address Assignment: MANUAL
Interface IPv6 oper status: Enabled
Link local IPv6 address: fe80::1618:77ff:fe25:6eb9/64
MTU 1532 bytes, IP MTU 1500 bytes
ARP type: ARPA, ARP Timeout: 60
Last clearing of "show interface" counters: 10:24:21
Queuing strategy: fifo
Input statistics:
    89 packets, 10056 octets
Output statistics:
    207 packets, 7376 octets
Time since last interface status change: 10:23:21
```

Supported releases 10.4.3.0 or later

show nve remote-vtep

Displays information about remote VXLAN tunnel endpoints.

Syntax `show nve remote-vtep [ip-address | summary | counters]`

Parameters

ip-address	Display detailed information about a specified remote VTEP.
summary	Display summary information about remote VTEPs.
counters	Display statistics on remote VTEP traffic.

Default Not configured

Command mode EXEC

Usage information Use the `show nve remote-vtep` command to display the IP address, operational state, and configured VXLANs for each remote VTEP. The remote MAC learning and unknown unicast drop settings used for each VXLAN ID (VNI) are also displayed.

Example

```
OS10# show nve remote-vtep summary
Remote-VTEP       State
-----------------------
2.2.2.2           up
```

```
OS10# show nve remote-vtep
Codes: DP - MAC-learn Dataplane, CP - MAC-learn Controlplane, UUD - Unknown-
Unicast-Drop
IP Address: 2.2.2.2,  State: up, Encap: VxLAN
   VNI list: 10000(DP), 200(DP), 300(DP)
```

Supported releases 10.4.2.0 or later

show nve remote-vtep counters

Displays VXLAN packet statistics for a remote VTEP.

Syntax `show nve remote-vtep [ip-address] counters`

Parameters

- *ip-address* — Enter IP address of a remote VTEP.

Default Not configured

Command mode EXEC

Usage information Use the `show nve remote-vtep counters` command to display input and output statistics for VXLAN traffic on a remote VTEP. A VTEP is identified by its IP address.

Use the `clear nve remote-vtep [ip-address] counters` command to clear VXLAN packet statistics.

Example

```
OS10# show nve remote-vtep counters
Peer               Input (Packets/Bytes)       Output (Packets/Bytes)
10.10.10.10         857/8570                    257/23709
20.20.20.20         457/3570                    277/13709
```

Supported releases 10.4.2.0 or later

show nve vxlan-vni

Displays information about the VXLAN virtual networks on the switch.

Syntax `show nve vxlan-vni`

Parameters None

Default Not configured

Command mode EXEC

Usage information Use the `show nve vxlan-vni` command to display information about configured VXLAN virtual networks. Each VXLAN virtual network is identified by its virtual-network ID.

Example

```
OS10# show nve vxlan-vni
VNI       Virtual-Network    Source-IP    Remote-VTEPs
-----------------------------------------------------
10000     1                  1.1.1.1      2.2.2.2
200       2                  1.1.1.1      2.2.2.2
300       300                1.1.1.1      2.2.2.2
```

Supported releases 10.4.2.0 or later

show virtual-network

Displays a virtual-network configuration, including all VXLAN configurations.

Syntax `show virtual-network [vn-id]`

Parameters

vn-id	Enter a virtual-network ID (1 to 65535).

Default Not configured

Command mode EXEC

Usage information Use the `show virtual-network` command to display the VNID, port members, source interface, and remote tunnel endpoints of a VXLAN virtual network.

Example

```
OS10# show virtual-network
Codes: DP - MAC-learn Dataplane, CP - MAC-learn Controlplane, UUD - Unknown-
Unicast-Drop
Un-tagged VLAN: 888
Virtual Network: 60000
   VLTi-VLAN: 2500
   Members:
      VLAN 1000: port-channel1, ethernet1/1/9, ethernet1/1/10
      VLAN 2500: port-channel1000
   VxLAN Virtual Network Identifier: 16775000
      Source Interface: loopback100(222.222.222.222)
      Remote-VTEPs (flood-list): 55.55.55.55(DP),77.1.1.1(DP)
```

Supported releases 10.4.2.0 or later

show virtual-network counters

Displays packet statistics for virtual networks.

Syntax `show virtual-network [vn-id] counters`

Parameters

vn-id	Enter a virtual-network ID (1 to 65535).

Default Not configured

Command mode EXEC

Usage information Use the `show virtual-network counters` command to monitor the packet throughput on virtual networks, including VXLANs.

Use the `clear virtual-network counters` command to clear virtual-network counters.

Example

```
OS10# show virtual-network counters
Virtual-Network       Input (Packets/Bytes)       Output (Packets/Bytes)
1000                  857/8570                    257/23709
2000                  457/3570                    277/13709
```

Supported releases 10.4.2.0 or later

show virtual-network interface counters

Displays packet statistics for a member port, port channel, or VLAN in VXLAN virtual networks.

Syntax `show virtual-network interface {ethernet node/slot/port:subport | port-channel number} [vlan vlan-id] counters`

Parameters

interface ethernet *node/slot/port[:subport]*	Enter the port information for an Ethernet interface.
interface port-channel *number*	Enter a port-channel number (1-128).
vlan *vlan-id*	(Optional) Enter a VLAN ID (1 to 4093).

Default Not configured

Command mode EXEC

Usage information Use the `show virtual-network interface counters` command to monitor the packet throughput on a port interface that is a member of a VXLAN virtual network. A VLAN member interface can be assigned to only one virtual network.

To clear VXLAN packet counters on a member port or VLAN members of a virtual network, enter the `clear virtual-network interface {ethernet node/slot/port:subport | port-channel number} [vlan vlan-id] counters` command.

Example

```
OS10# show virtual-network interface 1/1/3 vlan 100 counters
Virtual-Network      Input (Packets/Bytes)       Output (Packets/Bytes)
2000                 457/3570                    277/13709
```

Supported releases 10.4.2.0 or later

show virtual-network interface

Displays the VXLAN virtual networks and server VLANs to which a port is assigned.

Syntax `show virtual-network interface {ethernet node/slot/port:subport | port-channel number}`

Parameters

interface ethernet *node/slot/port[:subport]*	Enter the port information for an Ethernet interface.
interface port-channel *number*	Enter a port-channel number (1-128).

Default Not configured

Command mode EXEC

Usage information Use the `show virtual-network interface ethernet` command to verify the VXLAN VLANs in which an Ethernet port connected to downstream servers is a member.

Example

```
OS10# show virtual-network interface ethernet 1/1/1
Interface          Vlan      Virtual-network
ethernet1/1/1      100         1000
```

```
ethernet1/1/1       200          2000
ethernet1/1/1       300          3000
```

Supported releases 10.4.2.0 or later

show virtual-network vlan

Displays the VXLAN virtual networks to which a VLAN is assigned.

Syntax `show virtual-network vlan vlan-id`

Parameters

`vlan vlan-id` Enter a VLAN ID (1 to 4093).

Default Not configured

Command mode EXEC

Usage information Use the `show virtual-network vlan` command to verify the VXLAN virtual networks to which a VLAN is assigned, including the port members connected to downstream servers.

Example

```
OS10# show show virtual-network 100
Vlan  Virtual-network   Interface
100    1000             ethernet1/1/1,ethernet1/1/2
```

Supported releases 10.4.2.0 or later

show vlan (virtual network)

Displays the VLANs assigned to virtual networks.

Syntax `show vlan`

Parameters None

Default Not configured

Command mode EXEC

Usage information Use the `show vlan` command to display the VLAN port interfaces that transmit VXLAN packets over a virtual network.

Example

```
OS10# show vlan
Codes: * - Default VLAN, M - Management VLAN, R - Remote Port Mirroring VLANs,
@ - Attached to Virtual Network
Q: A - Access (Untagged), T - Tagged

  NUM  Status  Description  Q  Ports
* 1    up                   A  Eth1/1/1-1/1/48
@ 100  up                   T  Eth1/1/2,Eth1/1/3
                            A  Eth1/1/1
@ 101  up                   T  port-channel5
  200  up                   T  Eth1/1/11-1/1/15
```

Supported releases 10.4.2.0 or later

source-interface loopback

Configures a dedicated loopback interface as the source virtual tunnel endpoint (VTEP).

Syntax `source-interface loopback number`

Parameters

loopback *number*	Enter the loopback interface used as the source interface of a VXLAN virtual tunnel (0–16383).

Default Not configured

Command mode NVE-INSTANCE

Usage information The IP address of the loopback interface serves as the source IP address in encapsulated packets transmitted from the switch as an NVE VTEP. The `no` version of this command removes the configured value.

- The loopback interface must have an IP address configured. The loopback IP address must be reachable from the remote VTEP.
- You cannot change the source interface if at least one VXLAN virtual network ID (VNID) is configured for the NVE instance.

Enter the `source-interface loopback number` command in NVE mode to override a previously configured value and reconfigure the source IP address.

Examples

```
OS10(config-nve)# source-interface loopback 1
```

Supported releases 10.4.2.0 or later

virtual-network

Creates a virtual network for VXLAN tunneling.

Syntax `virtual-network vn-id`

Parameters

vn-id	Enter the virtual-network ID (1 to 65535).

Default Not configured

Command mode CONFIGURATION

Usage information The virtual network operates as a Layer 2 bridging domain. To add a VXLAN to the virtual network, use the `vxlan-vni` command. The `no` version of this command removes the configured virtual network.

Example

```
OS10(config)# virtual-network 1000
OS10(config-vn)#
```

Supported releases 10.4.2.0 or later

virtual-network untagged-vlan

Configures a dedicated VLAN for internal use to transmit untagged traffic on member ports in virtual networks on the switch.

Syntax `virtual-network untagged-vlan vlan-id`

Parameters

id	Enter the reserved untagged VLAN ID (1 to 4093).

Default Not configured

Command mode CONFIGURATION

Usage information The untagged VLAN ID is used internally for all untagged member interfaces that belong to virtual networks. You cannot use the reserved untagged VLAN ID for a simple VLAN bridge or for tagged traffic on member interfaces of virtual networks. The `no` version of this command removes the configured value.

Example

```
OS10(config)# virtual-network untagged-vlan 10
```

Supported releases 10.4.2.0 or later

vxlan-vni

Assigns a VXLAN ID to a virtual network.

Syntax `vxlan-vni vni`

Parameters

vni	Enter the VXLAN ID for a virtual network (1-16,777,215).

Default Not configured

Command mode VIRTUAL-NETWORK

Usage information The `vxlan-vni` associates a VXLAN ID number with a virtual network. The `no` version of this command removes the configured ID.

Example

```
OS10(conf-vn-100)# vxlan-vni 100
OS10(config-vn-vxlan-vni)#
```

Supported releases 10.4.2.0 or later

VXLAN MAC commands

clear mac address-table dynamic nve remote-vtep

Clears all MAC addresses learned from a remote VTEP.

Syntax `clear mac address-table dynamic nve remote-vtep ip-address`

Parameters

remote-vtep ***ip-address***	Clear MAC addresses learned from the specified remote VTEP.

Default Not configured

Command mode EXEC

Usage information To display the MAC addresses learned from a remote VTEP, use the `show mac address-table nve remote-vtep` command. Use the `clear mac address-table dynamic nve remote-vtep` command to delete all MAC address entries learned from a remote VTEP.

Example

```
OS10# clear mac address-table dynamic nve remote-vtep 32.1.1.1
```

Supported releases 10.4.2.0 or later

clear mac address-table dynamic virtual-network

Clears MAC addresses learned on all or a specified VXLAN virtual network.

Syntax

```
clear mac address-table dynamic virtual-network [interface {ethernet node/slot/
port:subport | port-channel number} | local | vn-id [address mac-address |
local]]
```

Parameters

interface ethernet *node/slot/port[:subport]*	Clear all MAC addresses learned on the specified interface.
interface port-channel *number*	Clear all MAC addresses learned on the specified port channel.
virtual-network *vn-id*	Clear all MAC addresses learned on the specified virtual network (1 to 65535).
local	Clear only locally-learned MAC addresses.
vn-id	Clear learned MAC addresses on the specified virtual network (1 to 65535).
***vn-id* local**	Clear locally learned MAC addresses on the specified virtual network (1 to 65535).
vn-id* address *mac-address	Clear only the MAC address entry learned in the specified virtual network. Enter the MAC address in *EEEE.EEEE.EEEE* format.

Default Not configured

Command mode EXEC

Usage information Use the `clear mac address-table dynamic virtual-network` command with no optional parameters to delete all dynamic MAC address entries that are learned only on virtual-network bridges from the MAC address table. This command does not delete MAC address entries learned on simple VLAN bridges. Use the `show mac address-table virtual-network` command to display the MAC addresses learned on a virtual network.

Example

```
OS10# clear mac address-table dynamic virtual-network
```

Supported releases 10.4.2.0 or later

show mac address-table count extended

Displays the number of MAC addresses learned on all VLANs and VXLAN virtual networks.

Syntax

```
show mac address-table count extended [interface {ethernet node/slot/
port:subport | port-channel number}]
```

Parameters

interface ethernet *node/slot/port[:subport]*	Display the number of MAC addresses learned on all VLANs and VXLANs on the specified interface.
interface port-channel *number*	Display the number of MAC addresses learned on all VLANs and VXLANs on the specified port channel.

Default Not configured

Command mode EXEC

Usage information Use the `show mac address-table count extended` command to display the number of MAC address entries learned on all VLANs and VXLAN virtual networks.

Example

```
OS10# show mac address-table count extended
MAC Entries for all vlans :
Dynamic Address Count :                  10
Static Address (User-defined) Count :    2
Total MAC Addresses in Use:              12
```

Supported releases 10.4.2.0 or later

show mac address-table count nve

Displays the number of MAC addresses learned on a VXLAN virtual network or from a remote VXLAN tunnel endpoint.

Syntax

```
show mac address-table count nve {vxlan-vni vni | remote-vtep ip-address}
```

Parameters

vxlan-vni *vni*	Display MAC addresses learned on the specified VXLAN virtual network (1-16,777,215).
remote-vtep *ip-address*	Display MAC addresses learned from the specified remote VTEP.

Default Not configured

Command mode EXEC

Usage information Use the `clear mac address-table dynamic nve remote-vtep` command to delete all MAC address entries learned from a remote VTEP. Use the `clear mac address-table dynamic virtual-network vn-id` command to delete all dynamic MAC address entries learned on a virtual-network bridge.

Example

```
OS10# show mac address-table count nve vxlan-vni 1001
MAC Entries for all vlans :
Dynamic Address Count :                  1
Static Address (User-defined) Count :    0
Total MAC Addresses in Use:              1
```

```
OS10# show mac address-table count nve remote-vtep 32.1.1.1
MAC Entries for all vlans :
Dynamic Address Count :                 2
Static Address (User-defined) Count :   0
Total MAC Addresses in Use:             2
```

Supported releases 10.4.2.0 or later

show mac address-table count virtual-network

Displays the number of MAC addresses learned on virtual networks.

Syntax `show mac address-table count virtual-network [dynamic | local | remote | static | interface {ethernet node/slot/port:subport | port-channel number} | vn-id]`

Parameters

dynamic	Display the number of local dynamically-learned MAC addresses.
local	Display the number of local MAC addresses.
remote	Display the number of MAC addresses learned from remote VTEPs.
static	Display the number of local statically-configured MAC addresses.
interface ethernet *node/slot/port[:subport]*	Display the number of MAC addresses learned on the specified interface.
interface port-channel *number*	Display the number of MAC addresses learned on the specified port channel.
vn-id	Display the number of MAC addresses learned on the specified virtual network (1-65535).

Default Not configured

Command mode EXEC

Usage information Use the `show mac address-table count virtual-network` command to display the number of MAC address entries learned on virtual networks in the MAC address table.

Example

```
OS10# show mac address-table count virtual-network
MAC Entries for all vlans :
Dynamic Address Count :                 8
Static Address (User-defined) Count :   0
Total MAC Addresses in Use:             8
```

Supported releases 10.4.2.0 or later

show mac address-table extended

Displays MAC addresses learned on all VLANs and VXLANs.

Syntax `show mac address-table extended [address mac-address | interface {ethernet node/slot/port:subport | port-channel number} | static | dynamic]`

Parameters

address *mac-address*	Display only information about the specified MAC address.
interface ethernet *node/slot/port[:subport]*	Display only MAC addresses learned on the specified interface.
interface port-channel *number*	Display only MAC addresses learned on the specified port channel.
static	Display only static MAC addresses.
dynamic	Display only dynamic MAC addresses.

Default Not configured

Command mode EXEC

Usage information By default, MAC learning from a remote VTEP is enabled. Use the `show mac address-table extended` command to verify the MAC addresses learned both on VXLAN virtual networks and VLANs on the switch. The `show mac address-table` command displays the MAC addresses learned only on LAN port and VLAN interfaces.

Example

```
OS10# show mac address-table extended
Virtual-Network  VlanId  MAC Address         Type      Interface/Remote-VTEP
-------------------------------------------------------------------------------
-                500     00:00:00:00:11:11   dynamic   ethernet1/1/31:1
-                500     00:00:00:00:44:44   dynamic   port-channel1000
-                1       aa:bb:cc:dd:f0:03   static    port-channel1000
-                500     aa:bb:cc:dd:f0:03   static    port-channel1000
-                4000    aa:bb:cc:dd:f0:03   static    port-channel1000
10000                    00:00:00:00:00:11   dynamic   ethernet1/1/31:1
10000            100     00:00:00:00:00:44   dynamic   port-channel1000
10000            100     00:00:00:00:00:55   dynamic   port-channel10
10000                    00:00:00:00:00:77   dynamic   VxLAN(32.1.1.1)
20000            300     00:00:00:00:00:22   dynamic   port-channel100
20000            300     00:00:00:00:00:33   dynamic   port-channel1000
20000            300     00:00:00:00:00:66   dynamic   port-channel10
20000                    00:00:00:00:00:88   dynamic   VxLAN(32.1.1.1)
```

Supported releases 10.4.2.0 or later

show mac address-table nve

Displays the MAC addresses learned on a VXLAN virtual network or from a remote VXLAN tunnel endpoint.

Syntax `show mac address-table nve {vxlan-vni` *vni* `| remote-vtep` *ip-address*`}`

Parameters

vxlan-vni *vni*	Display MAC addresses learned on the specified VXLAN virtual network (1-16,777,215).
remote-vtep *ip-address*	Display MAC addresses learned from the specified remote VTEP.

Default Not configured

Command mode EXEC

Usage information Use the `clear mac address-table dynamic nve remote-vtep` command to delete all MAC address entries learned from a remote VTEP. Use the `clear mac address-table dynamic virtual-network` *`vn-id`* command to delete all dynamic MAC address entries learned on a virtual-network bridge.

Example

```
OS10# show mac address-table nve remote-vtep 32.1.1.1
Virtual-Network  VNI    MAC Address        Type     Remote-VTEP
-----------------------------------------------------------------
10000            9999   00:00:00:00:00:77  dynamic  VxLAN(32.1.1.1)
20000            19999  00:00:00:00:00:88  dynamic  VxLAN(32.1.1.1)

OS10# show mac address-table nve vxlan-vni 9999
Virtual-Network  VNI    MAC Address        Type     Remote-VTEP
-----------------------------------------------------------------
10000            9999   00:00:00:00:00:77  dynamic  VxLAN(32.1.1.1)
```

Supported releases 10.4.2.0 or later

show mac address-table virtual-network

Displays the MAC addresses learned on all or a specified virtual network.

Syntax `show mac address-table virtual-network [`*`vn-id`* `| local | remote | static | dynamic | address` *`mac-address`* `| interface {ethernet` *`node/slot/port:subport`* `| port-channel` *`number`*`}]`

Parameters

vn-id	Display only information about the specified virtual network.
local	Display only locally-learned MAC addresses.
remote	Display only remote MAC addresses.
static	Display only static MAC addresses.
dynamic	Display only dynamic MAC addresses.
address *mac-address*	Display only information about the specified MAC address. Enter the MAC address in *EEEE.EEEE.EEEE* format.
interface ethernet *node/slot/port[:subport]*	Display only MAC addresses learned on the specified interface.
interface port-channel *number*	Display only MAC addresses learned on the specified port channel.

Default Not configured

Command mode EXEC

Usage information By default, MAC learning from a remote VTEP is enabled. Use the `show mac address-table virtual-network` command to verify the MAC addresses learned on VXLAN virtual networks.

Example

```
OS10# show mac address-table virtual-network
Virtual-Network     VlanId       MAC Address         Type        Interface/Remote-VTEP
-------------------------------------------------------------------------------------
10000                            00:00:00:00:00:11   dynamic     ethernet1/1/31:1
10000               100          00:00:00:00:00:44   dynamic     port-channel1000
10000               100          00:00:00:00:00:55   dynamic     port-channel10
10000                            00:00:00:00:00:77   dynamic     VxLAN(32.1.1.1)
```

```
10000                100          34:a0:a0:a1:a2:f6   dynamic     port-channel10
20000                300          00:00:00:00:00:22   dynamic     port-channel100
20000                300          00:00:00:00:00:33   dynamic     port-channel1000
20000                300          00:00:00:00:00:66   dynamic     port-channel10
20000                             00:00:00:00:00:88   dynamic     VxLAN(32.1.1.1)
20000                300          34:a0:a0:a1:a2:f6   dynamic     port-channel10
```

Supported releases 10.4.2.0 or later

Example: VXLAN with static VTEP

This example uses a typical Clos leaf-spine topology with static VXLAN tunnel endpoints (VTEPs) in VLT dual-homing domains. Individual switch configuration shows how to set up an end-to-end VXLAN. The underlay IP network routes are advertised using OSPF.

- On VTEPs 1 and 2, access ports are assigned to the virtual network using a switch-scoped VLAN configuration.
- On VTEPs 3 and 4, access ports are assigned to the virtual network using a port-scoped VLAN configuration.
- Overlay routing between hosts in different IP subnets is configured on the VTEPs.

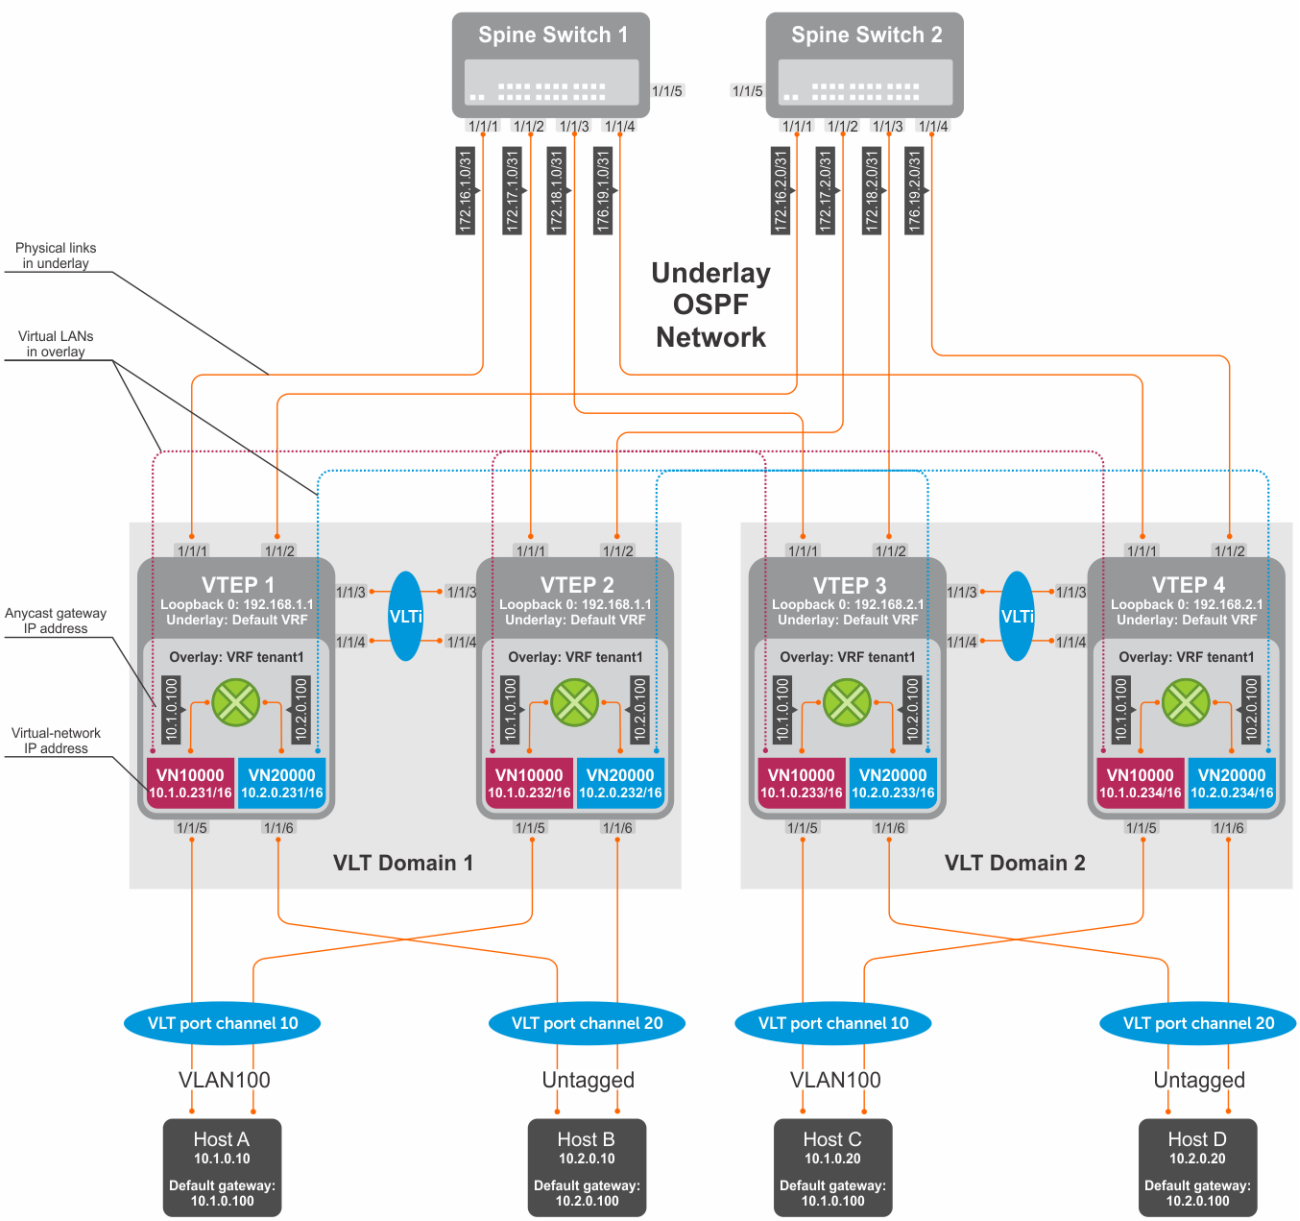


Figure 6. Static VXLAN use case

VTEP 1 Leaf Switch

1. Configure the underlay OSPF protocol

Do not configure the same IP address for the router ID and the source loopback interface in Step 2.

```
OS10(config)# router ospf 1
OS10(config-router-ospf-1)# router-id 172.16.0.1
OS10(config-router-ospf-1)# exit
```

2. Configure a loopback interface

```
OS10(config)# interface loopback0
OS10(conf-if-lo-0)# no shutdown
OS10(conf-if-lo-0)# ip address 192.168.1.1/32
OS10(conf-if-lo-0)# ip ospf 1 area 0.0.0.0
OS10(conf-if-lo-0)# exit
```

3. Configure the loopback interface as the VXLAN source tunnel interface

```
OS10(config)# nve
OS10(config-nve)# source-interface loopback0
OS10(config-nve)# exit
```

4. Configure VXLAN virtual networks with a static VTEP

```
OS10(config)# virtual-network 10000
OS10(config-vn-10000)# vxlan-vni 10000
OS10(config-vn-vxlan-vni)# remote-vtep 192.168.2.1
OS10(config-vn-vxlan-vni-remote-vtep)# exit
OS10(config-vn-vxlan-vni)# exit
OS10(config-vn-10000)# exit
OS10(config)# virtual-network 20000
OS10(config-vn-20000)# vxlan-vni 20000
OS10(config-vn-vxlan-vni)# remote-vtep 192.168.2.1
OS10(config-vn-vxlan-vni-remote-vtep)# exit
OS10(config-vn-vxlan-vni)# exit
OS10(config-vn-20000)# exit
```

5. Assign VLAN member interfaces to virtual networks

Use either a switch-scoped VLAN-to-VNI mapping:

```
OS10(config)# interface vlan100
OS10(config-if-vl-100)# virtual-network 10000
OS10(config-if-vl-100)# no shutdown
OS10(config-if-vl-100)# exit
OS10(config)# interface vlan200
OS10(config-if-vl-100)# virtual-network 20000
OS10(config-if-vl-100)# no shutdown
OS10(config-if-vl-100)# exit
```

6. Configure access ports as VLAN members for a switch-scoped VLAN-to-VNI mapping

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# no shutdown
OS10(conf-if-po-10)# switchport mode trunk
OS10(conf-if-po-10)# switchport trunk allowed vlan 100
OS10(conf-if-po-10)# exit

OS10(config)# interface ethernet1/1/5
OS10(conf-if-eth1/1/5)# no shutdown
OS10(conf-if-eth1/1/5)# channel-group 10 mode active
OS10(conf-if-eth1/1/5)# no switchport
OS10(conf-if-eth1/1/5)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# no shutdown
OS10(conf-if-po-20)# switchport access vlan 200
OS10(conf-if-po-20)# exit

OS10(config)# interface ethernet1/1/6
OS10(conf-if-eth1/1/6)# no shutdown
OS10(conf-if-eth1/1/6)# channel-group 20 mode active
OS10(conf-if-eth1/1/6)# no switchport
OS10(conf-if-eth1/1/6)# exit
```

7. Configure upstream network-facing ports

```
OS10(config)# interface ethernet1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/1)# ip address 172.16.1.0/31
OS10(conf-if-eth1/1/1)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/1)# exit

OS10(config)# interface ethernet1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/2)# ip address 172.16.2.0/31
OS10(conf-if-eth1/1/2)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/2)# exit
```

8. Configure VLT

Configure dedicated L3 underlay path to reach VLT Peer in case of network failure

```
OS10(config)# interface vlan4000
OS10(config-if-vl-4000)# no shutdown
OS10(config-if-vl-4000)# ip address 172.16.250.1/30
OS10(config-if-vl-4000)# ip ospf 1 area 0.0.0.0
OS10(config-if-vl-4000)# exit
```

Configure VLT port channel

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# vlt-port-channel 10
OS10(conf-if-po-10)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# vlt-port-channel 20
OS10(conf-if-po-20)# exit
```

Configure VLTi member links

```
OS10(config)# interface ethernet1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# exit

OS10(config)# interface ethernet1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# exit
```

Configure VLT domain

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination 10.16.150.1
OS10(conf-vlt-1)# discovery-interface ethernet1/1/3,1/1/4
OS10(conf-vlt-1)# vlt-mac aa:bb:cc:dd:ee:ff
OS10(conf-vlt-1)# exit
```

Configure UFD with uplink VLT ports and downlink network ports

```
OS10(config)# uplink-state-group 1
OS10(conf-uplink-state-group-1)# enable
OS10(conf-uplink-state-group-1)# downstream ethernet1/1/1-1/1/2
OS10(conf-uplink-state-group-1)# upstream port-channel10
OS10(conf-uplink-state-group-1)# upstream port-channel20
OS10(conf-uplink-state-group-1)# exit
```

9. Configure overlay IP routing

Create tenant VRF

```
OS10(config)# ip vrf tenant1
OS10(conf-vrf)# exit
```

Configure anycast L3 gateway MAC address for all VTEPs

```
OS10(config)# ip virtual-router mac-address 00:01:01:01:01:01
```

Configure routing with anycast gateway IP address for each virtual network

```
OS10(config)# interface virtual-network 10000
OS10(config-if-vn-10000)# ip vrf forwarding tenant1
OS10(config-if-vn-10000)# ip address 10.1.0.231/16
OS10(config-if-vn-10000)# ip virtual-router address 10.1.0.100
OS10(config-if-vn-10000)# no shutdown
OS10(config-if-vn-10000)# exit
OS10(config)# interface virtual-network 20000
OS10(config-if-vn-20000)# ip vrf forwarding tenant1
OS10(config-if-vn-20000)# ip address 10.2.0.231/16
OS10(config-if-vn-20000)# ip virtual-router address 10.2.0.100
OS10(config-if-vn-20000)# no shutdown
OS10(config-if-vn-20000)# exit
```

VTEP 2 Leaf Switch

1. Configure the underlay OSPF protocol

Do not configure the same router ID on other VTEP switches.

```
OS10(config)# router ospf 1
OS10(config-router-ospf-1)# router-id 172.17.0.1
OS10(config-router-ospf-1)# exit
```

2. Configure a loopback interface

The source-interface IP address must be same as the source-interface IP address on the VLT peer.

```
OS10(config)# interface loopback0
OS10(conf-if-lo-0)# no shutdown
OS10(conf-if-lo-0)# ip address 192.168.1.1/32
OS10(conf-if-lo-0)# ip ospf 1 area 0.0.0.0
OS10(conf-if-lo-0)# exit
```

3. Configure the loopback interface as the VXLAN source tunnel interface

```
OS10(config)# nve
OS10(config-nve)# source-interface loopback0
OS10(config-nve)# exit
```

4. Configure VXLAN virtual networks with a static VTEP

```
OS10(config)# virtual-network 10000
OS10(config-vn-10000)# vxlan-vni 10000
OS10(config-vn-vxlan-vni)# remote-vtep 192.168.2.1
OS10(config-vn-vxlan-vni-remote-vtep)# exit
OS10(config-vn-vxlan-vni)# exit
OS10(config-vn-10000)# exit
OS10(config)# virtual-network 20000
OS10(config-vn-20000)# vxlan-vni 20000
OS10(config-vn-vxlan-vni)# remote-vtep 192.168.2.1
OS10(config-vn-vxlan-vni-remote-vtep)# exit
```

```
OS10(config-vn-vxlan-vni)# exit
OS10(config-vn-20000)# exit
```

5. Assign switch-scoped VLAN to a virtual network

```
OS10(config)# interface vlan100
OS10(config-if-vl-100)# virtual-network 10000
OS10(config-if-vl-100)# no shutdown
OS10(config-if-vl-100)# exit
OS10(config)# interface vlan200
OS10(config-if-vl-100)# virtual-network 20000
OS10(config-if-vl-100)# no shutdown
OS10(config-if-vl-100)# exit
```

6. Configure access ports as VLAN members

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# no shutdown
OS10(conf-if-po-10)# switchport mode access
OS10(conf-if-po-10)# switchport access vlan 200
OS10(conf-if-po-10)# exit

OS10(config)# interface ethernet1/1/5
OS10(conf-if-eth1/1/5)# no shutdown
OS10(conf-if-eth1/1/5)# channel-group 10 mode active
OS10(conf-if-eth1/1/5)# no switchport
OS10(conf-if-eth1/1/5)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# no shutdown
OS10(conf-if-po-20)# switchport mode access
OS10(conf-if-po-20)# switchport access vlan 200
OS10(conf-if-po-20)# exit

OS10(config)# interface ethernet1/1/6
OS10(conf-if-eth1/1/6)# no shutdown
OS10(conf-if-eth1/1/6)# channel-group 20 mode active
OS10(conf-if-eth1/1/6)# no switchport
OS10(conf-if-eth1/1/6)# exit
```

7. Configure upstream network-facing ports

```
OS10(config)# interface ethernet1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/1)# ip address 172.17.1.0/31
OS10(conf-if-eth1/1/1)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/1)# exit

OS10(config)# interface ethernet1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/2)# ip address 172.17.2.0/31
OS10(conf-if-eth1/1/2)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/2)# exit
```

8. Configure VLT

Configure dedicated L3 underlay path to reach VLT Peer in case of network failure

```
OS10(config)# interface vlan4000
OS10(config-if-vl-4000)# no shutdown
OS10(config-if-vl-4000)# ip address 172.16.250.2/30
OS10(config-if-vl-4000)# ip ospf 1 area 0.0.0.0
OS10(config-if-vl-4000)# exit
```

Configure VLT port channel

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# vlt port-channel 10
OS10(conf-if-po-10)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# vlt port-channel 20
OS10(conf-if-po-20)# exit
```

Configure VLTi member links

```
OS10(config)# interface ethernet1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# exit

OS10(config)# interface ethernet1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# exit
```

Configure VLT domain

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination 10.16.150.2
OS10(conf-vlt-1)# discovery-interface ethernet1/1/3,1/1/4
OS10(conf-vlt-1)# vlt-mac aa:bb:cc:dd:ee:ff
OS10(conf-vlt-1)# exit
```

Configure UFD with uplink VLT ports and downlink network ports

```
OS10(config)# uplink-state-group 1
OS10(conf-uplink-state-group-1)# enable
OS10(conf-uplink-state-group-1)# downstream ethernet1/1/1-1/1/2
OS10(conf-uplink-state-group-1)# upstream port-channel10
OS10(conf-uplink-state-group-1)# upstream port-channel20
OS10(conf-uplink-state-group-1)# exit
```

9. Configure overlay IP routing

Create tenant VRF

```
OS10(config)# ip vrf tenant1
OS10(conf-vrf)# exit
```

Configure anycast L3 gateway MAC address for all VTEPs

```
OS10(config)# ip virtual-router mac-address 00:01:01:01:01:01
```

Configure routing with anycast gateway IP address for each virtual network

```
OS10(config)# interface virtual-network 10000
OS10(config-if-vn-10000)# ip vrf forwarding tenant1
OS10(config-if-vn-10000)# ip address 10.1.0.232/16
OS10(config-if-vn-10000)# ip virtual-router address 10.1.0.100
OS10(config-if-vn-10000)# no shutdown
OS10(config-if-vn-10000)# exit
OS10(config)# interface virtual-network 20000
OS10(config-if-vn-20000)# ip vrf forwarding tenant1
OS10(config-if-vn-20000)# ip address 10.2.0.232/16
OS10(config-if-vn-20000)# ip virtual-router address 10.2.0.100
OS10(config-if-vn-20000)# no shutdown
OS10(config-if-vn-20000)# exit
```

VTEP 3 Leaf Switch

1. Configure the underlay OSPF protocol

Do not configure the same IP address for the router ID and the source loopback interface in Step 2.

```
OS10(config)# router ospf 1
OS10(config-router-ospf-1)# router-id 172.18.0.1
OS10(config-router-ospf-1)# exit
```

2. Configure a loopback interface

```
OS10(config)# interface loopback0
OS10(conf-if-lo-0)# no shutdown
OS10(conf-if-lo-0)# ip address 192.168.2.1/32
OS10(conf-if-lo-0)# ip ospf 1 area 0.0.0.0
OS10(conf-if-lo-0)# exit
```

3. Configure the loopback interface as the VXLAN source tunnel interface

```
OS10(config)# nve
OS10(config-nve)# source-interface loopback0
OS10(config-nve)# exit
```

4. Configure VXLAN virtual networks with a static VTEP

```
OS10(config)# virtual-network 10000
OS10(config-vn-10000)# vxlan-vni 10000
OS10(config-vn-vxlan-vni)# remote-vtep 192.168.1.1
OS10(config-vn-vxlan-vni-remote-vtep)# exit
OS10(config-vn-vxlan-vni)# exit
OS10(config-vn-10000)# exit
OS10(config)# virtual-network 20000
OS10(config-vn-20000)# vxlan-vni 20000
OS10(config-vn-vxlan-vni)# remote-vtep 192.168.1.1
OS10(config-vn-vxlan-vni-remote-vtep)# exit
OS10(config-vn-vxlan-vni)# exit
OS10(config-vn-20000)# exit
```

5. Configure reserved VLAN ID for untagged member interfaces

```
OS10(config)# virtual-network untagged-vlan 1000
```

6. Configure access ports

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# no shutdown
OS10(conf-if-po-10)# switchport mode trunk
OS10(conf-if-po-10)# no switchport access vlan
OS10(conf-if-po-10)# exit

OS10(config)# interface ethernet1/1/5
OS10(conf-if-eth1/1/5)# no shutdown
OS10(conf-if-eth1/1/5)# channel-group 10 mode active
OS10(conf-if-eth1/1/5)# no switchport
OS10(conf-if-eth1/1/5)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# no shutdown
OS10(conf-if-po-20)# switchport mode trunk
OS10(conf-if-po-20)# no switchport access vlan
OS10(conf-if-po-20)# exit

OS10(config)# interface ethernet1/1/6
OS10(conf-if-eth1/1/6)# no shutdown
```

```
OS10(conf-if-eth1/1/6)# channel-group 20 mode active
OS10(conf-if-eth1/1/6)# no switchport
OS10(conf-if-eth1/1/6)# exit
```

7. Add access ports to VXLAN virtual networks

```
OS10(config)# virtual-network 10000
OS10(config-vn-10000)# member-interface port-channel 10 vlan-tag 100
OS10(config-vn-10000)# exit
OS10(config)# virtual-network 20000
OS10(config-vn-20000)# member-interface port-channel 20 untagged
OS10(config-vn-20000)# exit
```

NOTE: This step shows how to add access ports using a port-scoped VLAN-to-VNI mapping. You can also add access ports using a switch-scoped VLAN-to-VNI mapping. However, you cannot use both methods at the same time; you must use either a port-scoped or switch-scoped VLAN-to-VNI mapping.

8. Configure upstream network-facing ports

```
OS10(config)# interface ethernet1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/1)# ip address 172.18.1.0/31
OS10(conf-if-eth1/1/1)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/1)# exit

OS10(config)# interface ethernet1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/2)# ip address 172.18.2.0/31
OS10(conf-if-eth1/1/2)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/2)# exit
```

9. Configure VLT

Configure VLTi VLAN for VXLAN virtual network

```
OS10(config)# virtual-network 10000
OS10(config-vn-10000)# vlti-vlan 100
OS10(config-vn-10000)# exit
OS10(config)# virtual-network 20000
OS10(config-vn-20000)# vlti-vlan 200
OS10(config-vn-20000)# exit
```

Configure dedicated L3 underlay path to reach VLT Peer in case of network failure

```
OS10(config)# interface vlan4000
OS10(config-if-vl-4000)# no shutdown
OS10(config-if-vl-4000)# ip address 172.16.250.9/30
OS10(config-if-vl-4000)# ip ospf 1 area 0.0.0.0
OS10(config-if-vl-4000)# exit
```

Configure VLT port channel

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# vlt port-channel 10
OS10(conf-if-po-10)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# vlt port-channel 20
OS10(conf-if-po-20)# exit
```

Configure VLTi member links

```
OS10(config)# interface ethernet1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
```

```
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# exit

OS10(config)# interface ethernet1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# exit
```

Configure VLT domain

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination 10.16.150.3
OS10(conf-vlt-1)# discovery-interface ethernet1/1/3,1/1/4
OS10(conf-vlt-1)# vlt-mac aa:bb:dd:cc:ff:ee
OS10(conf-vlt-1)# exit
```

Configure UFD with uplink VLT ports and downlink network ports

```
OS10(config)# uplink-state-group 1
OS10(conf-uplink-state-group-1)# enable
OS10(conf-uplink-state-group-1)# downstream ethernet1/1/1-1/1/2
OS10(conf-uplink-state-group-1)# upstream port-channel10
OS10(conf-uplink-state-group-1)# upstream port-channel20
OS10(conf-uplink-state-group-1)# exit
```

10. Configure overlay IP routing

Create tenant VRF

```
OS10(config)# ip vrf tenant1
OS10(conf-vrf)# exit
```

Configure anycast L3 gateway

```
OS10(config)# ip virtual-router mac-address 00:01:01:01:01:01
```

Configure routing with anycast gateway IP address for each virtual network

```
OS10(config)# interface virtual-network 10000
OS10(config-if-vn-10000)# ip vrf forwarding tenant1
OS10(config-if-vn-10000)# ip address 10.1.0.233/16
OS10(config-if-vn-10000)# ip virtual-router address 10.1.0.100
OS10(config-if-vn-10000)# no shutdown
OS10(config-if-vn-10000)# exit
OS10(config)# interface virtual-network 20000
OS10(config-if-vn-20000)# ip vrf forwarding tenant1
OS10(config-if-vn-20000)# ip address 10.2.0.233/16
OS10(config-if-vn-20000)# ip virtual-router address 10.2.0.100
OS10(config-if-vn-20000)# no shutdown
OS10(config-if-vn-20000)# exit
```

VTEP 4 Leaf Switch

1. Configure the underlay OSPF protocol

Do not configure the same IP address for the router ID and the source loopback interface in Step 2.

```
OS10(config)# router ospf 1
OS10(config-router-ospf-1)# router-id 172.19.0.1
OS10(config-router-ospf-1)# exit
```

2. Configure a loopback interface

```
OS10(config)# interface loopback0
OS10(conf-if-lo-0)# no shutdown
```

```
OS10(conf-if-lo-0)# ip address 192.168.2.1/32
OS10(conf-if-lo-0)# ip ospf 1 area 0.0.0.0
OS10(conf-if-lo-0)# exit
```

3. Configure the loopback interface as the VXLAN source tunnel interface

```
OS10(config)# nve
OS10(config-nve)# source-interface loopback0
OS10(config-nve)# exit
```

4. Configure VXLAN virtual networks with a static VTEP

```
OS10(config)# virtual-network 10000
OS10(config-vn-10000)# vxlan-vni 10000
OS10(config-vn-vxlan-vni)# remote-vtep 192.168.1.1
OS10(config-vn-vxlan-vni-remote-vtep)# exit
OS10(config-vn-vxlan-vni)# exit
OS10(config-vn-10000)# exit
OS10(config)# virtual-network 20000
OS10(config-vn-20000)# vxlan-vni 20000
OS10(config-vn-vxlan-vni)# remote-vtep 192.168.1.1
OS10(config-vn-vxlan-vni-remote-vtep)# exit
OS10(config-vn-vxlan-vni)# exit
OS10(config-vn-20000)# exit
```

5. Configure reserved VLAN ID for untagged member interfaces

```
OS10(config)# virtual-network untagged-vlan 1000
```

6. Configure access ports

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# no shutdown
OS10(conf-if-po-10)# switchport mode trunk
OS10(conf-if-po-10)# no switchport access vlan
OS10(conf-if-po-10)# exit

OS10(config)# interface ethernet1/1/5
OS10(conf-if-eth1/1/5)# no shutdown
OS10(conf-if-eth1/1/5)# channel-group 10 mode active
OS10(conf-if-eth1/1/5)# no switchport
OS10(conf-if-eth1/1/5)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# no shutdown
OS10(conf-if-po-20)# switchport mode trunk
OS10(conf-if-po-20)# no switchport access vlan
OS10(conf-if-po-20)# exit

OS10(config)# interface ethernet1/1/6
OS10(conf-if-eth1/1/6)# no shutdown
OS10(conf-if-eth1/1/6)# channel-group 20 mode active
OS10(conf-if-eth1/1/6)# no switchport
OS10(conf-if-eth1/1/6)# exit
```

7. Add access ports to VXLAN virtual network

```
OS10(config)# virtual-network 10000
OS10(config-vn-10000)# member-interface port-channel 10 vlan-tag 100
OS10(config-vn-10000)# exit
OS10(config)# virtual-network 20000
OS10(config-vn-20000)# member-interface port-channel 20 untagged
OS10(config-vn-20000)# exit
```

8. Configure upstream network-facing ports

```
OS10(config)# interface ethernet1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
```

```
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/1)# ip address 172.19.1.0/31
OS10(conf-if-eth1/1/1)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/1)# exit

OS10(config)# interface ethernet1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/2)# ip address 172.19.2.0/31
OS10(conf-if-eth1/1/2)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/2)# exit
```

9. Configure VLT

Configure VLTi VLAN for VXLAN virtual network

```
OS10(config)# virtual-network 10000
OS10(config-vn-10000)# vlti-vlan 200
OS10(config-vn-10000)# exit
OS10(config)# virtual-network 20000
OS10(config-vn-20000)# vlti-vlan 100
OS10(config-vn-20000)# exit
```

Configure dedicated L3 underlay path to reach VLT Peer in case of network failure

```
OS10(config)# interface vlan4000
OS10(config-if-vl-4000)# no shutdown
OS10(config-if-vl-4000)# ip address 172.16.250.10/30
OS10(config-if-vl-4000)# ip ospf 1 area 0.0.0.0
OS10(config-if-vl-4000)# exit
```

Configure VLT port channel

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# vlt port-channel 10
OS10(conf-if-po-10)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# vlt port-channel 20
OS10(conf-if-po-20)# exit
```

Configure VLTi member links

```
OS10(config)# interface ethernet1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# exit

OS10(config)# interface ethernet1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# exit
```

Configure VLT domain

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination 10.16.150.4
OS10(conf-vlt-1)# discovery-interface ethernet1/1/3,1/1/4
OS10(conf-vlt-1)# vlt-mac aa:bb:dd:cc:ff:ee
OS10(conf-vlt-1)# exit
```

Configure UFD with uplink VLT ports and downlink network ports

```
OS10(config)# uplink-state-group 1
OS10(conf-uplink-state-group-1)# enable
OS10(conf-uplink-state-group-1)# downstream ethernet1/1/1-1/1/2
OS10(conf-uplink-state-group-1)# upstream port-channel10
```

```
OS10(conf-uplink-state-group-1)# upstream port-channel20
OS10(conf-uplink-state-group-1)# exit
```

10. Configure overlay IP routing

Create tenant VRF

```
OS10(config)# ip vrf tenant1
OS10(conf-vrf)# exit
```

Configure anycast L3 gateway for all VTEPs in all virtual networks

```
OS10(config)# ip virtual-router mac-address 00:01:01:01:01:01
```

Configure routing with anycast gateway IP address for each virtual network

```
OS10(config)# interface virtual-network 10000
OS10(config-if-vn-10000)# ip vrf forwarding tenant1
OS10(config-if-vn-10000)# ip address 10.1.0.234/16
OS10(config-if-vn-10000)# ip virtual-router address 10.1.0.100
OS10(config-if-vn-10000)# no shutdown
OS10(config-if-vn-10000)# exit
OS10(config)# interface virtual-network 20000
OS10(config-if-vn-20000)# ip vrf forwarding tenant1
OS10(config-if-vn-20000)# ip address 10.2.0.234/16
OS10(config-if-vn-20000)# ip virtual-router address 10.2.0.100
OS10(config-if-vn-20000)# no shutdown
OS10(config-if-vn-20000)# exit
```

Spine Switch 1

1. Configure downstream ports on underlay links to leaf switches

```
OS10(config)# interface ethernet1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# ip address 172.16.1.1/31
OS10(conf-if-eth1/1/1)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/1)# exit

OS10(config)# interface ethernet1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/2)# ip address 172.17.1.1/31
OS10(conf-if-eth1/1/2)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/2)# exit

OS10(config)# interface ethernet1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# ip address 172.18.1.1/31
OS10(conf-if-eth1/1/3)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/3)# exit

OS10(config)# interface ethernet1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# ip address 172.19.1.1/31
OS10(conf-if-eth1/1/4)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/4)# exit
```

2. Configure the underlay OSPF protocol

```
OS10(config)# router ospf 1
OS10(config-router-ospf-1)# router-id 172.200.0.1
OS10(config-router-ospf-1)# exit
```

Spine Switch 2

1. Configure downstream ports on underlay links to leaf switches

```
OS10(config)# interface ethernet1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# ip address 172.16.2.1/31
OS10(conf-if-eth1/1/1)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/1)# exit

OS10(config)# interface ethernet1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/2)# ip address 172.17.2.1/31
OS10(conf-if-eth1/1/2)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/2)# exit

OS10(config)# interface ethernet1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# ip address 172.18.2.1/31
OS10(conf-if-eth1/1/3)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/3)# exit

OS10(config)# interface ethernet1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# ip address 172.19.2.1/31
OS10(conf-if-eth1/1/4)# ip ospf 1 area 0.0.0.0
OS10(conf-if-eth1/1/4)# exit
```

2. Configure the underlay OSPF protocol

```
OS10(config)# router ospf 1
OS10(config-router-ospf-1)# router-id 172.201.0.1
OS10(config-router-ospf-1)# exit
```

BGP EVPN for VXLAN

Ethernet Virtual Private Network (EVPN) is a control plane for VXLAN that is used to reduce flooding in the network and resolve scalability concerns. EVPN uses MP-BGP to exchange information between VTEPs. EVPN was introduced in RFC 7432 and is based on BGP MPLS-based VPNs. RFC 8365 describes VXLAN-based EVPN.

The MP-BGP EVPN control plane provides protocol-based remote VTEP discovery, and MAC and ARP learning. As a result, flooding related to L2 unknown unicast traffic is reduced. The distribution of host MAC and IP reachability information supports virtual machine mobility and scalable VXLAN overlay network designs.

The BGP EVPN protocol groups MAC addresses and ARP/neighbor addresses under EVPN instances (EVIs) in order to exchange them between VTEPs. In OS10, each EVI is associated with a VXLAN VNI in a 1:1 mapping.

Benefits of BGP EVPN-based VXLAN

- Eliminates the flood-and-learn method of VTEP discovery by enabling control-plane learning of end-host L2 and L3 reachability information.
- Minimizes the network flooding of unknown unicast and broadcast traffic through EVPN-based MAC and IP route advertisements on local VTEPs.
- Provides support for host mobility.

BGP EVPN compared to static VXLAN

OS10 supports two types of VXLAN NVO overlay networks:

- Static VXLAN
- BGP EVPN

Static VXLAN and BGP EVPN for VXLAN are configured and operate in the same ways:

- The overlay and underlay networks are manually configured.
- Each virtual network and VNI are manually configured.
- Access port membership in a virtual network is manually configured.
- Underlay reachability to VTEP peers is provisioned or learned using existing routing protocols.

Static VXLAN and BGP EVPN for VXLAN configuration and operation differ as described in the following table.

Table 22. Differences between Static VXLAN and VXLAN BGP EVPN

Static VXLAN	VXLAN BGP EVPN
To start sending and receiving virtual-network traffic to and from a remote VTEP, you must manually configure the VTEP as a member of the virtual network.	No manual configuration is required. Each remote VTEP is automatically learned as a member of a virtual network from the EVPN routes received from the remote VTEP. After a remote VTEP's address is learned, VXLAN traffic is sent to, and received from, the VTEP.
Remote hosts are learned from data packets after decapsulation of the VXLAN header in the data plane.	Remote host MAC addresses are learned in the control plane using BGP EVPN Type 2 routes and MAC/IP advertisements.

VXLAN BGP EVPN operation

The EVPN address family allows VXLAN to carry EVPN routes in eBGP and iBGP sessions. In a data center network, you can use eBGP or iBGP for route exchange in both the IP underlay network and EVPN.

This sample BGP EVPN topology shows a typical data center (leaf-spine) network in which eBGP is used for exchanging IP routes in the IP underlay network, and EVPN routes in the VXLAN overlay network. All spine nodes are in one autonomous system (AS 65535). All leaf nodes are in another autonomous system (AS 65000).

To advertise underlay IP routes, eBGP peer sessions are established between the leaf and spine nodes using an interface IP address. To advertise EVPN routes, eBGP peer sessions between the leaf and spine nodes use a loopback IP address.

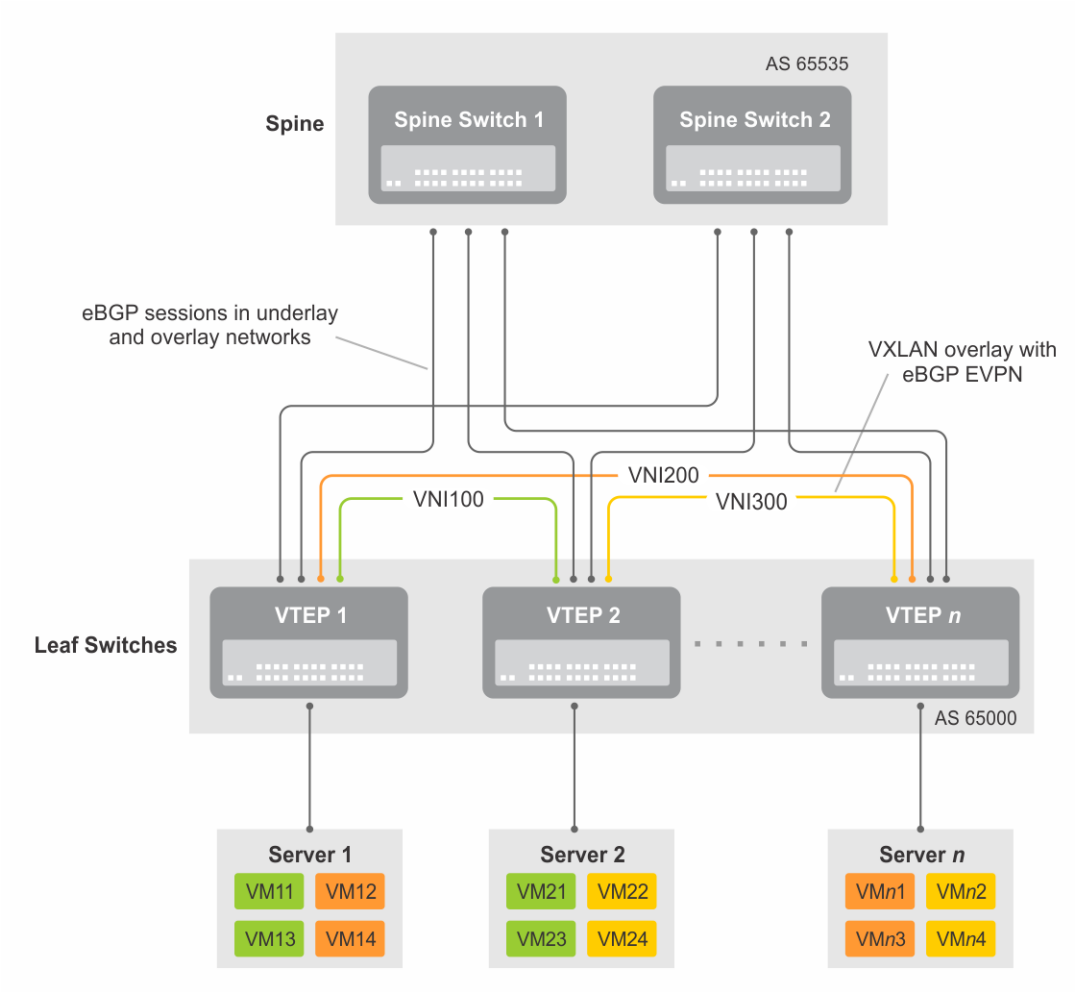


Figure 7. BGP EVPN topology

Leaf nodes

The leaf nodes are typically top-of-rack (TOR) switches in a data center network. They act as the VXLAN tunnel endpoints and perform VXLAN encapsulation and decapsulation. The leaf nodes also participate in the MP-BGP EVPN to support control plane and data plane functions.

The control plane functions include:

- Initiate and maintain route adjacencies using any routing protocol in the underlay network.
- Advertise locally learned routes to all MP-BGP EVPN peers.
- Process the routes received from remote MP-BGP EVPN peers and install them in the local forwarding plane.

The data plane functions include:

- Encapsulate server traffic with VXLAN headers and forward the packets in the underlay network.
- Decapsulate VXLAN packets received from remote VTEPs and forward the native packets to downstream hosts.
- Perform underlay route processing, including routing based on the outer IP address.

Spine nodes

The role of a spine node changes based on its control plane and data plane functions. Spine nodes participate in underlay route processing to forward packets and in the overlay network to advertise EVPN routes to all MP-BGP peers.

The control plane functions include:

- Initiate BGP peering with all neighbor leaf nodes.
- Advertise BGP routes to all BGP peers.
- In the underlay network, initiate and maintain the routing adjacencies with all leaf and spine nodes.

The data plane functions include:

- Perform only the underlay route processing based on the outer header in VXLAN encapsulated packets.
- Does not perform VXLAN encapsulation or decapsulation.

BGP EVPN running on each VTEP listens to the exchange of route information in the local overlay, encodes the learned routes as BGP EVPN routes, and injects them into BGP to be advertised to the peers. Tunnel endpoints are advertised as Type 3 EVPN routes. MAC/IP addresses are advertised as Type 2 EVPN routes.

EVPN instance

An EVPN instance (EVI) spans across the VTEPs that participate in an Ethernet VPN. Each virtual network (tenant segment) that needs to be advertised using EVPN must be associated with an EVI. In OS10, you can configure EVIs in auto-EVI or manual configuration mode.

- Auto-EVI — After you configure a virtual network on a VTEP, auto-EVI mode automatically creates an EVPN instance. The route distinguisher and route target values are automatically generated:
  - The EVI ID is auto-generated with the same value as the virtual-network ID (VNID) configured on the VTEP and is associated with the VXLAN network ID (VNI).
  - A Route Distinguisher (RD) is auto-generated for each EVI ID. A Route Distinguisher maintains the uniqueness of an EVPN route between different EVPN instances.
  - A Route Target (RT) import and export value are auto-generated for each EVI ID. A Route Target determines how EVPN routes are distributed among EVPN instances.
- Manual EVI configuration — If you need to specify the Route Distinguisher and Route Target values, manually configure EVPN instances and associate each EVI with the overlay virtual network using the VXLAN VNI. The EVI is activated only when the virtual network, Router Distinguisher, and Route Target values are configured.

  In manual EVI configuration, you can either manually configure the RD and RT or have them auto-configured.

Route distinguisher

The Route Distinguisher is an 8-byte identifier, which uniquely identifies an EVPN instance (EVI). Each EVPN route is prefixed with a unique RD and exchanged between BGP peers, making the tenant's route unique across the network. In this way, overlapping address spaces among tenants are supported.

A Route Distinguisher for each EVI can be auto-generated or manually configured. In auto-EVI mode, the RD is auto-generated. In manual EVI configuration mode, the RD can be auto-generated or manually configured.

As specified in RFC 7432, a manually configured RD are encoded in the format: *4-octet-ipv4-address*:*2-octet-number*. An auto-generated RD has the format: *vtep-ip-address*:*evi*.

Route target

While a Route Distinguisher maintains the uniqueness of an EVPN route among different EVIs, a Route Target controls the way in which EVPN routes are distributed among EVIs. Each EVI is configured with an import and export Route Target value. BGP EVPN routes advertised for an EVI carry the export Route Target associated with the EVI. A receiving VTEP downloads information in the BGP EVPN route to EVIs that have a matching import Route Target value.

The Route Target import and export for each EVI can be auto-generated or manually configured. In auto-EVI mode, the RT is auto-generated. In manual EVI configuration mode, the RT can be auto-generated or manually configured.

The Route Target consists of a 2-octet *type* and a 6-octet *value*. If a Route Target is auto-configured, the encoding format is different for a 2-byte and 4-byte AS number (ASN):

- For a 2-byte ASN, the RT *type* is set to 0200 (Type 0 in RFC 4364). The RT *value* is encoded in the format described in section 5.1.2.1 of RFC 8365: *2-octet-ASN*: *4-octet-number*, where the following values are used in the *4-octet-number* field:
  - Type: 1
  - D-ID: 0
  - Service-ID: VNI
- For a 4-byte ASN, OS10 can auto-configure RTs for both 2-byte and 4-byte ASNs. The RT *type* is set to 0202 (Type 2 in RFC 4364). The RT *value* is encoded in the format: *4-octet-ASN*: *2-octet-number*, where the *2-octet-number* field contains the EVI ID. In auto-EVI mode, the EVI ID is the same as the virtual network ID (VNID). Therefore, in 4-byte ASN deployment, OS10 supports RT auto-configuration if the VNID-to-VNI mapping is the same on all VTEPs.

Configure BGP EVPN for VXLAN

To set up BGP EVPN service in a VXLAN overlay network:

1 Configure the VXLAN overlay network (see Configure VXLAN). If you enable routing for VXLAN virtual networks, integrated routing and bridging (IRB) for BGP EVPN is automatically enabled.
2 Configure BGP to advertise EVPN routes.
3 Configure EVPN, including the VNI, RD, and RT values associated with the EVPN instance.
4 Verify the BGP EVPN configuration.

Usage guidelines

- Only EVPN bridging (L2 gateway) functionality is supported.
- Only EVPN route types 2 and 3 are supported.
- OS10 only supports asymmetric IRB. Symmetric IRB will be supported in future releases.

Configuration

1 Configure BGP to advertise EVPN routes.

EVPN requires that you establish MP-BGP sessions between leaf and spine nodes in the underlay network. On each spine and leaf node, configure at least two BGP peering sessions:

- A directly connected BGP peer in the underlay network to advertise VTEP and loopback IP addresses using the IPv4 unicast address family
- A BGP peer in the overlay network to advertise overlay information using the EVPN address family. In BGP peer sessions in the overlay, activate only the EVPN address family.

For each BGP peer session in the underlay network:

a Create a BGP instance in CONFIGURATION mode. You enter router BGP configuration mode.

```
router bgp as-number
```

b Assign an IP address to the BGP instance in ROUTER-BGP mode.

```
router-id ip-address
```

c Enter IPv4 address-family configuration mode from ROUTER-BGP mode.

```
address-family ipv4 unicast
```

d Advertise the IPv4 prefix to BGP peers in the address family in ROUTER-BGP-ADDRESS-FAMILY mode.

```
network ip-address/mask
```

e Return to ROUTER-BGP mode.

```
exit
```

f Configure the BGP peer address in ROUTER-BGP mode.

```
neighbor ip-address
```

g Assign the BGP neighbor to an autonomous system in ROUTER-BGP-NEIGHBOR mode.

```
remote-as as-number
```

h Enable the peer session with the BGP neighbor in ROUTER-BGP-NEIGHBOR mode.

```
no shutdown
```

i Return to ROUTER-BGP mode.

```
exit
```

For each BGP peer session in the overlay network:

a Configure the BGP peer using its loopback IP address on the VTEP in ROUTER-BGP mode.

```
neighbor loopback-ip-address
```

b Assign the BGP neighbor loopback address to the autonomous system in ROUTER-BGP-NEIGHBOR mode. The neighbor loopback IP address is used as the source interface on the remote VTEP.

```
remote-as as-number
```

c Use the local loopback address as the source address in BGP packets sent to the neighbor in ROUTER-BGP-NEIGHBOR mode.

```
update-source loopback0
```

d Send an extended community attribute to the BGP neighbor in ROUTER-BGP-NEIGHBOR mode.

```
send-community extended
```

e Enable the peer session with the BGP neighbor in ROUTER-BGP-NEIGHBOR mode.

```
no shutdown
```

f Configure the L2 VPN EVPN address family for VXLAN host-based routing to the BGP peer in ROUTER-BGP-NEIGHBOR mode.

```
address-family l2vpn evpn
```

g Enable the exchange of L2VPN EVPN addresses with the BGP peer in ROUTER-BGP-NEIGHBOR mode.

```
activate
```

h Return to ROUTER-BGP mode.

```
exit
```

i Enter IPv4 address-family configuration mode from ROUTER-BGP mode.

```
address-family ipv4 unicast
```

j Disable the exchange of IPv4 addresses with BGP peers in ROUTER-BGP mode.

```
no activate
```

k Return to ROUTER-BGP-NEIGHBOR mode.

```
exit
```

l (Optional) If all the leaf switches are configured in the same ASN:

- On each leaf switch, enter L2VPN EVPN address-family configuration mode from ROUTER-BGP-NEIGHBOR mode. Activate the exchange of L2VPN EVPN addresses with BGP peers. Configure the switch so that it can accept a route with the local AS number in updates received from a peer in ROUTER-BGP-NEIGHBOR-AF mode. For example:

```
OS10(config-router-bgp-neighbor)# address-family l2vpn evpn
OS10(config-router-neighbor-af)# activate
OS10(config-router-neighbor-af)# allowas-in 1
OS10(config-router-neighbor-af)# exit
OS10(config-router-bgp-neighbor)# exit
```

- On each spine switch, disable sender-side loop detection to leaf switch neighbors in ROUTER-BGP-NEIGHBOR-AF mode.

```
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# exit
```

m (Optional) In a VLT deployment, on each leaf switch, configure the number of multi-hop peer routes in ROUTER-BGP-NEIGHBOR mode to ensure that the BGP EVPN peer session is established over the VLT VTEP peer if all local links to spine switches are down. For example:

```
OS10(conf-router-neighbor)# ebgp-multihop 1
```

2 Configure EVPN.

An EVPN instance (EVI) spans across the VTEPs that participate in the EVPN. In OS10, you can configure an EVI in auto-EVI or manual configuration mode.

- **Auto-EVI mode**

  1 Enable the EVPN control plane in CONFIGURATION mode.

  ```
  evpn
  ```

2 Enable auto-EVI creation for overlay virtual networks in EVPN mode. Auto-EVI creation is supported only if BGP EVPN is used with 2-byte AS numbers and if at least one BGP instance is enabled with the EVPN address family. No further manual configuration is allowed in auto-EVI mode.

```
auto-evi
```

- **Manual EVI configuration mode**

1 Enable the EVPN control plane in CONFIGURATION mode.

```
evpn
```

2 Manually create an EVPN instance in EVPN mode (1 to 65535).

```
evi id
```

3 Configure the Route Distinguisher in EVPN EVI mode.

```
rd {A.B.C.D:[1-65535] | auto}
```

Where:

- `rd A.B.C.D:[1-65535]` configures the Route Distinguisher with a 4-octet IPv4 address followed by a 2-octet-number.
- `rd auto` automatically generates the Route Distinguisher.

4 Configure the Route Target values in EVPN EVI mode.

```
route-target {auto | value [asn4] {import | export | both}}
```

Where:

- `route-target auto` auto-configures an import and export value for EVPN routes.
- `route-target value [asn4]{import | export | both}` configures an import or export value for EVPN routes in the format *2-octet-ASN*:*4-octet-number* or *4-octet-ASN*:*2-octet-number*.
  - The *2-octet* ASN number is 1 to 65535.
  - The *4-octet* ASN number is 1 to 4294967295.

  To configure the same value for the RT import and export values, use the `both` option. `asn4` advertises a 2-byte AS number as a 4-byte route target value. If you specify the `asn4` option, configure the VXLAN network ID associated with the EVPN instance in EVPN EVI mode (1 to 16,777,215). You must configure the same VNI value that you configured for the VXLAN virtual network (see Configure VXLAN).

  ```
  vni vni
  ```

3 Verify the BGP EVPN configuration.

Display EVPN instance configuration

```
OS10# show evpn evi 1
EVI : 65447, State : up
  Bridge-Domain       : (Virtual-Network)100, (VNI)100
  Route-Distinguisher : 1:110.111.170.102:65447(auto)
  Route-Targets       : 0:101:268435556(auto) both
  Inclusive Multicast : 110.111.170.107
```

Display VXLAN overlay for EVPN instance

```
OS10# show evpn vxlan-vni
VXLAN-VNI  EVI  Virtual-Network-Instance
100001     1    1
100010     2    2
```

Display BGP neighbors in EVPN instances

```
OS10# show ip bgp neighbors 110.111.170.102
BGP neighbor is 110.111.170.102, remote AS 100, local AS 100 internal link
BGP version 4, remote router ID 110.111.170.102
BGP state ESTABLISHED, in this state for 04:02:59
Last read 00:21:21 seconds
Hold time is 180, keepalive interval is 60 seconds
Configured hold time is 180, keepalive interval is 60 seconds
Fall-over disabled

Received 311 messages
```

```
  2 opens, 2 notifications, 3 updates
  304 keepalives, 0 route refresh requests
Sent 307 messages
  4 opens, 0 notifications, 2 updates
  301 keepalives, 0 route refresh requests
Minimum time between advertisement runs is 30 seconds
Minimum time before advertisements start is 0 seconds
Capabilities received from neighbor for IPv4 Unicast:
  MULTIPROTO_EXT(1)
  ROUTE_REFRESH(2)
  CISCO_ROUTE_REFRESH(128)
  4_OCTET_AS(65)
  MP_L2VPN_EVPN
Capabilities advertised to neighbor for IPv4 Unicast:
  MULTIPROTO_EXT(1)
  ROUTE_REFRESH(2)
  CISCO_ROUTE_REFRESH(128)
  4_OCTET_AS(65)
  MP_L2VPN_EVPN
Prefixes accepted 1, Prefixes advertised 1
Connections established 2; dropped 0
Last reset never
Prefixes ignored due to:
  Martian address 0, Our own AS in AS-PATH 0
  Invalid Nexthop 0, Invalid AS-PATH length 0
  Wellknown community 0, Locally originated 0

Local host: 110.111.180.195, Local port: 43081
Foreign host: 110.111.170.102, Foreign port: 179
```

Display BGP L2VPN EVPN address family

```
OS10# show ip bgp l2vpn evpn
BGP local RIB : Routes to be Added , Replaced , Withdrawn
BGP local router ID is 110.111.170.102
Status codes: s suppressed, S stale, d dampened, h history, * valid, > best
Path source: I - internal, a - aggregate, c - confed-external,
r - redistributed/network, S - stale
Origin codes: i - IGP, e - EGP, ? - incomplete
          Network                                 Next Hop         Metric  LocPrf  Weight  Path
*>r       Route distinguisher: 110.111.170.102:65447
[3]:[0]:[32]:[110.111.170.102]/152                110.111.170.102  0       100     32768   ?
*>        Route distinguisher: 110.111.170.107:64536
[3]:[0]:[32]:[110.111.170.107]/152                110.111.170.107  0       100     0       100
101 ?
```

Display EVPN routes for host MAC addresses

```
OS10# show evpn mac
Type  -(lcl): Local (rmt): remote

EVI     Mac-Address        Type      Seq-No  Interface/Next-Hop
50      00:00:00:aa:aa:aa  rmt       0       55.1.1.3
50      00:00:00:cc:cc:cc  lcl       0       ethernet1/1/8:1

OS10# show evpn mac evi 50
Type  -(lcl): Local (rmt): remote

EVI     Mac-Address        Type      Seq-No  Interface/Next-Hop
50      00:00:00:aa:aa:aa  rmt       0       55.1.1.3
50      00:00:00:cc:cc:cc  lcl       0       ethernet1/1/8:1
```

VXLAN BGP EVPN routing

Configure BGP EVPN for VXLAN describes how EVPN facilitates traffic switching within the same Layer 2 tenant segment (virtual network) on a VTEP for virtual networks that are associated with EVPN instances (EVIs). This section describes how EVPN implements overlay routing between L2 segments associated with EVIs belonging to the same tenant on a VTEP. The IETF draft draft-ietf-bess-evpn-inter-subnet-forwarding-05 describes EVPN inter-subnet forwarding — Integrated Routing and Bridging (IRB) — and how to use EVPN with IP routing between L2 tenant domains.

As described in Configure VXLAN — Enable overlay routing between virtual networks, you set up overlay routing by assigning a VRF to each tenant, creating a virtual-network interface, and assigning an IP subnet in the VRF to each virtual-network interface. The VTEP acts as the L3 gateway that routes traffic from one tenant subnet to another in the overlay before encapsulating it in the VXLAN header and transporting it over the underlay fabric. On virtual networks that are associated with EVIs, EVPN integrated routing and bridging (IRB) is enabled only after you create a virtual-network interface.

When IRB is enabled for a virtual network/EVI, EVPN operation on each VTEP also advertises the local tenant IP-MAC bindings learned on the EVPN-enabled virtual networks to all other VTEPs. The local tenant IP-MAC bindings are learned from ARP or ICMPv6 protocol operation. They are advertised as EVPN Type-2 BGP route updates to other VTEPs, each of whom then imports and installs them as ARP/IPv6 neighbor entries in the dataplane.

To enable the most efficient traffic forwarding on a VTEP, OS10 supports distributed gateway routing. A distributed gateway allows multiple VTEPs to act as the gateway router for a tenant subnet. The VTEP that is located nearest to a host acts as its gateway router.

To enable L3 gateway/IRB functionality for BGP EVPN, configure a VXLAN overlay network and enable routing on a switch, including:

1 Create a non-default VRF instance for overlay routing. For multi-tenancy, create a VRF instance for each tenant.
2 Globally configure the anycast gateway MAC address used by all VTEPs.
3 Configure a virtual-network interface for each virtual network, (optional) assign it to the tenant VRF, and configure an IP address. Then enable the interface.
4 Configure an anycast gateway IP address for each virtual network. OS10 supports distributed gateway routing.

For detailed information, see Configure VXLAN — Enable overlay routing between virtual networks.

EVPN supports different types of IRB routing for tenants — VMs and servers — that are connected to each VTEP in a tenant network.

- Asymmetric routing: IP routing is performed on ingress VTEPs. L2 bridging is performed on egress VTEPs. An ingress VTEP must be configured with a virtual network even for destination IP subnets that have no locally attached hosts. EVPN asymmetric IRB installs ARP entries to associate each tenant VM IP address with its corresponding VM MAC address in the overlay.
- Symmetric routing: IP routing is performed on both ingress and egress VTEPs. An ingress VTEP does not have to be configured with a virtual network for destination IP subnets that have no locally attached hosts. EVPN symmetric IRB installs ARP entries to associate each tenant VM destination IP address with the MAC address of the VTEP where the VM is located, reducing the number of required hardware next-hop routing resources.

NOTE: In release 10.4.3.0, OS10 supports only distributed asymmetric routing mode.

For a sample BGP EVPN routing configuration, see Example: VXLAN with BGP EVPN.

BGP EVPN with VLT

OS10 supports BGP EVPN operation between VLT peers that are configured as VTEPs. The configurations and best practices to set up VLT for VXLAN as described in the Configure VXLAN — Configure VLT also apply to BGP EVPN for VXLAN.

A recommended best practice is to configure iBGP peering for the IPv4 address family between the VTEPs in a VLT pair on a dedicated L3 VLAN that is used when connectivity to the underlay L3 network is lost. It is NOT required to enable the EVPN address family on the iBGP peering session between the VTEPs in a VLT pair because EVPN peering to the spine switch is performed on loopback interfaces.

Both VTEPs in a VLT pair advertise identical EVPN routes, which provides redundancy in case one of the VTEP peers fails. To set up redundant EVPN route advertisement, configure the same EVI, RD, and RT values for each VXLAN ID (VNI) on both VTEPs in a VLT pair, including:

- In auto-EVI mode, this identical configuration is automatically ensured if the virtual-network ID (VNID)-to-VNI association is the same on both VTEP peers.
- In manual EVI mode, you must configure the same EVI-to-VNID association on both VTEP peers.
- In manual EVI mode, you must configure the same RD and RT values on both VTEP peers.

In an EVPN configuration, increase the VLT delay-restore timer to allow for BGP EVPN adjacency to be established and for the remote MAC and neighbor entries to be downloaded by EVPN and installed in the data-plane. The VLT delay-restore determines the amount of time the VLT LAGs are kept operationally down at bootup to allow the data plane to set up and forward traffic, resulting in minimal traffic loss as the VLT peer node boots up and joins the VLT domain.

For a sample BGP EVPN VLT configuration, see Example: VXLAN with BGP EVPN.

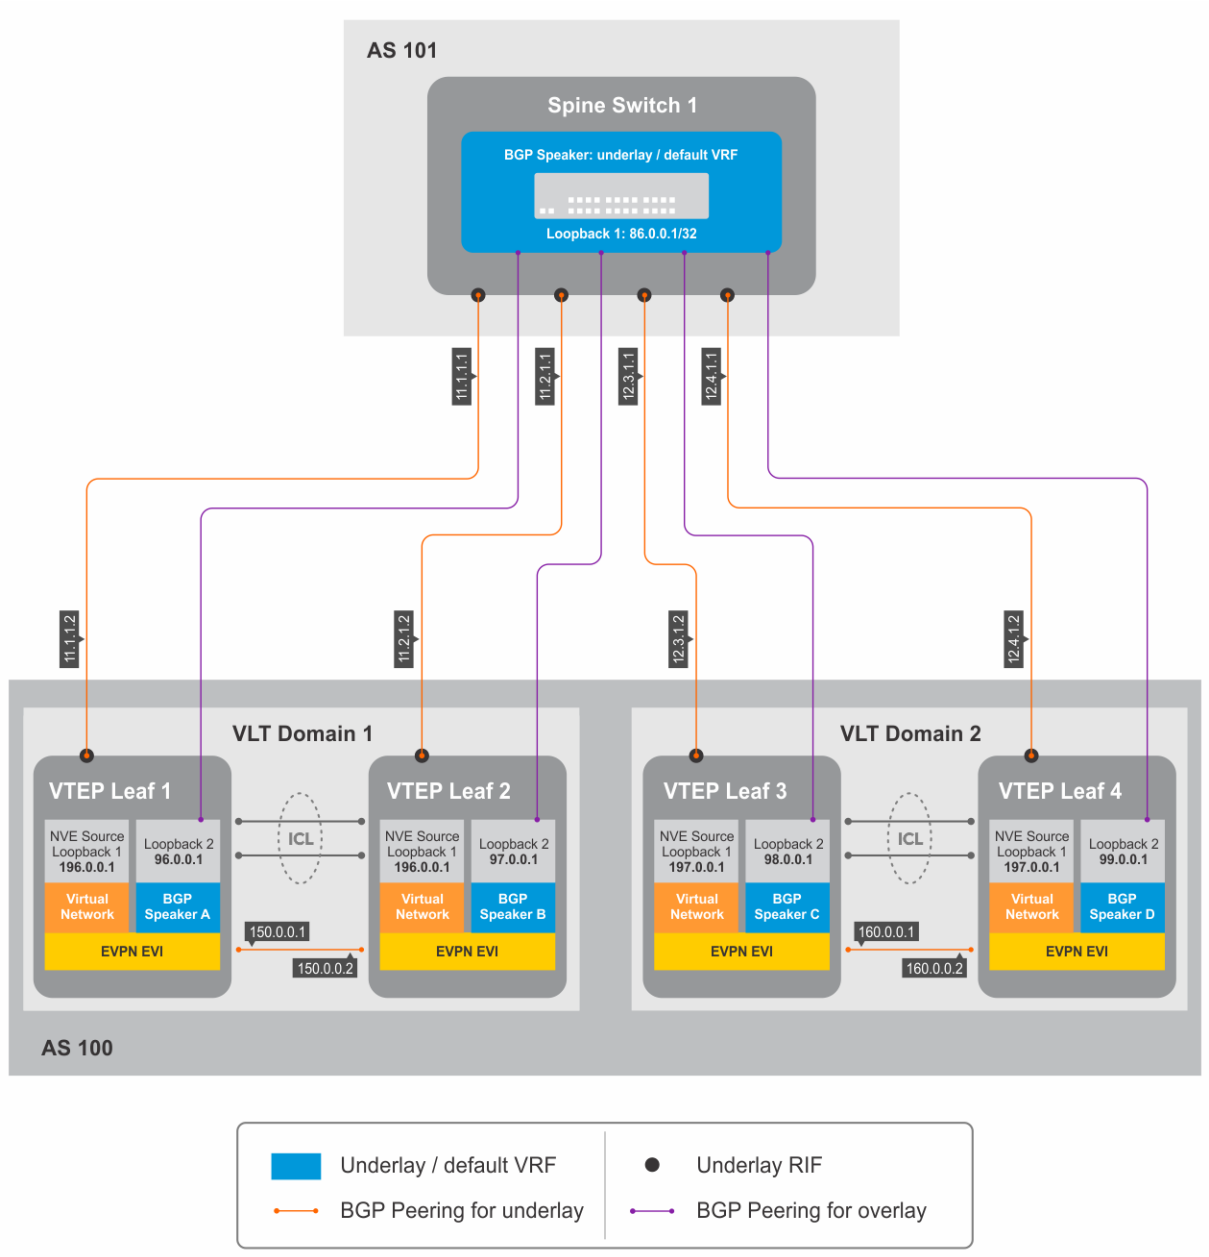


Figure 8. BGP EVPN in VLT domain

VXLAN BGP commands

activate (l2vpn evpn)

Enables the exchange of L2VPN EVPN address family information with a BGP neighbor or peer group.

Syntax	`activate`
Parameters	None

Default	Not configured
Command Mode	ROUTER-BGP-NEIGHBOR-AF
Usage Information	Use the `activate` command to exchange L2 VPN EVPN address information for VXLAN host-based routing with a BGP neighbor. The IPv4 unicast address family is enabled by default. To enable an address family exchange with a BGP neighbor, use the `activate` command. To disable an address family with a neighbor, use the `no activate` command.
Example	`OS10(conf-router-neighbor)# address-family l2vpn evpn unicast` `OS10(conf-router-bgp-neighbor-af)# activate`
Supported Releases	10.2.0E or later

address-family l2vpn evpn

Configures the L2 VPN EVPN address family for VXLAN host-based routing to a BGP neighbor.

Syntax	`address-family l2vpn evpn`
Parameters	None
Default	Not configured
Command mode	ROUTER-NEIGHBOR
Usage information	To use BGP EVPN service in a VXLAN, you must configure and enable the L2VPN EVPN address family on a VTEP to support host-based routing to each BGP neighbor.

Example	`OS10(config)# router bgp 100`
	`OS10(config-router-bgp-100)# neighbor 45.0.0.1`
	`OS10(config-router-neighbor)# address-family l2vpn evpn`

Supported releases	10.4.2.0 or later

allowas-in

Configures the number of times the local AS number can appear in the BGP AS_PATH path attribute before the switch rejects the route.

Syntax	`allowas-in` *`as-number`*
Parameters	*`as-number`*—Enter the number of occurrences for a local AS number, from 1 to 10.
Default	Disabled
Command Mode	ROUTER-BPG-NEIGHBOR-AF
Usage Information	Use this command to enable the BGP speaker to accept a route with the local AS number in updates received from a peer for the specified number of times. The `no` version of this command resets the value to the default.
Example (IPv4)	`OS10(config-router-neighbor)# address-family ipv4 unicast` `OS10(conf-router-bgp-neighbor-af)# allowas-in 5`
Example (IPv6)	`OS10(conf-router-template)# address-family ipv6 unicast` `OS10(conf-router-bgp-template-af)# allowas-in 5`
Supported Releases	10.3.0E or later

sender-side-loop-detection

Enables the sender-side loop detection process for a BGP neighbor.

Syntax `sender-side-loop-detection`

Parameters None

Default Enabled

Command Mode ROUTER-BGP-NEIGHBOR-AF

Usage Information This command helps detect routing loops, based on the AS path before it starts advertising routes. To configure a neighbor to accept routes use the `neighbor allowas-in` command. The `no` version of this command disables sender-side loop detection for that neighbor.

Example (IPv4)

```
OS10(conf-router-bgp-102)# neighbor 3.3.3.1
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-bgp-neighbor-af)# sender-side-loop-detection
```

Example (IPv6)

```
OS10(conf-router-bgp-102)# neighbor 32::1
OS10(conf-router-neighbor)# address-family ipv6 unicast
OS10(conf-router-bgp-neighbor-af)# no sender-side-loop-detection
```

Supported Releases 10.3.0E or later

show ip bgp l2vpn evpn

Displays the internal BGP routes in the L2VPN EVPN address family in EVPN instances.

Syntax `show ip bgp l2vpn evpn [summary | neighbors]`

Parameters

summary	Display a summary of the BGP routes in the L2VPN address family that are exchanged with remote VTEPs.
neighbors	Display the remote VTEPs with whom BGP routes in the L2VPN address family are exchanged.

Default Not configured

Command mode EXEC

Usage information Use the `show ip bgp l2vpn evpn` command to display the BGP routes used for the L2VPN EVPN address family in EVPN instances on the switch.

Examples

```
OS10# show ip bgp l2vpn evpn
BGP local RIB : Routes to be Added , Replaced , Withdrawn
BGP local router ID is 110.111.170.102
Status codes: s suppressed, S stale, d dampened, h history, * valid, > best
Path source: I - internal, a - aggregate, c - confed-external,
r - redistributed/network, S - stale
Origin codes: i - IGP, e - EGP, ? - incomplete
          Network                     Next Hop         Metric  LocPrf
Weight  Path
*>r       Route distinguisher: 110.111.170.102:65447
[3]:[0]:[32]:[110.111.170.102]/152      110.111.170.102  0       100
32768   ?
*>        Route distinguisher: 110.111.170.107:64536
```

```
[3]:[0]:[32]:[110.111.170.107]/152      110.111.170.107  0       100
0       100 101 ?
```

```
OS10# show ip bgp l2vpn evpn summary
BGP router identifier 2.2.2.2 local AS number 4294967295
Neighbor    AS          MsgRcvd     MsgSent     Up/Down           State/Pfx
3.3.3.3     4294967295  2831        9130        05:57:27          504
4.4.4.4     4294967295  2364        9586        05:56:43          504
5.5.5.5     4294967295  4947        8399        01:10:39          11514
6.6.6.6     4294967295  2413        7310        05:51:56          504
```

```
OS10# show ip bgp l2vpn evpn neighbors
BGP neighbor is 3.3.3.3, remote AS 4294967295, local AS 4294967295  internal
link

  BGP version 4, remote router ID 3.3.3.3
  BGP state ESTABLISHED, in this state for 06:21:55
  Last read 00:37:43 seconds
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time is 180, keepalive interval is 60 seconds
  Fall-over disabled
  Route reflector client

  Received 2860 messages
    1 opens, 0 notifications, 2422 updates
    437 keepalives, 0 route refresh requests
  Sent 32996 messages
    1 opens, 0 notifications, 32565 updates
    430 keepalives, 0 route refresh requests
 Minimum time between advertisement runs is 30 seconds
  Minimum time before advertisements start is 0 seconds

  Capabilities received from neighbor for IPv4 Unicast:
    ROUTE_REFRESH(2)
    CISCO_ROUTE_REFRESH(128)
    4_OCTET_AS(65)
    MP_L2VPN_EVPN(1)
  Capabilities advertised to neighbor for IPv4 Unicast:
    ROUTE_REFRESH(2)
    CISCO_ROUTE_REFRESH(128)
    4_OCTET_AS(65)
    MP_L2VPN_EVPN(1)
  Prefixes accepted 504, Prefixes advertised 13012
  Connections established 1; dropped 0
  Last reset never
  Local host: 2.2.2.2, Local port: 37853
  Foreign host: 3.3.3.3, Foreign port: 179
  ...
```

Supported releases 10.4.2.0 or later

VXLAN EVPN commands

auto-evi

Automatically creates an EVPN instance, including Route Distinguisher and Route Target values.

Syntax `auto-evi`

Parameters None

Default Not configured

Command mode EVPN

Usage information In deployments running BGP with 2-byte or 4-byte autonomous systems, auto-EVI automatically creates EVPN instances when you create a virtual network on a VTEP in the overlay network. In auto-EVI mode, the Route Distinguisher and Route Target values are automatically generated:

- For a 2-byte autonomous system:
  - The RD auto-configures as Type 1 from the overlay network source IP address and the auto-generated EVI index.
  - The RT auto-configures as Type 0 from the 2-byte AS and the 3-byte VNI (Type encoded as 0x0002).
- For a 4-byte autonomous system:
  - The RD auto-configures as Type 1 from the overlay network source IP address and the auto-generated EVI index.
  - The RT auto-configures as Type 2 from the 4-byte AS and the 2-byte EVI (Type encoded as 0x0202).

Example

```
OS10(config)# evpn
OS10(config-evpn)# auto-evi
```

Supported releases 10.4.2.0 or later

evi

Creates an EVPN instance (EVI) in EVPN mode.

Syntax `evi id`

Parameters

id	Enter the EVPN instance ID (1 to 65535).

Default Not configured

Command mode EVPN

Usage information If an MP-BGP network uses 4-byte autonomous systems or if you need to specify the Route Distinguisher and Route Target values, manually configure EVPN instances and associate each EVI with the overlay VXLAN virtual network. The EVI is activated only when the VXLAN network ID (VNI), Router Distinguisher, Route Target, and virtual network are configured.

Example

```
OS10(config)# evpn
OS10(config-evpn)# evi 10
OS10(config-evpn-evi)#
```

Supported releases 10.4.2.0 or later

evpn

Enable the EVPN control plane for VXLAN.

Syntax `evpn`

Parameters None

Default Not configured

Command mode CONFIGURATION

Usage information Enabling EVPN triggers BGP to advertise EVPN capability with AFI=25 and SAFI=70 to all BGP peers in an autonomous system. The `no evpn` command disables EVPN on the switch.

Example

```
OS10(config)# evpn
OS10(config-evpn)#
```

Supported releases 10.4.2.0 or later

rd

Configures the Route Distinguisher (RD) value used in EVPN routes.

Syntax `rd {A.B.C.D:[1-65535] | auto}`

Parameters

A.B.C.D: **[1-65535]**	Manually configure the route distinguisher with a 4-octet IPv4 address followed by a 2-octet-number. (1-65535).
auto	Configure the route distinguisher to be automatically generated.

Default Not configured

Command mode EVPN-EVI

Usage information A route distinguisher (RD) maintains the uniqueness of an EVPN route between different EVPN instances. The RD auto-configures as Type 1 from the overlay network source IP address and the auto-generated EVPN instance ID.

Example

```
OS10(config)# evpn
OS10(config-evpn)# evi 10
OS10(config-evpn-evi)# vni 10000
OS10(config-evpn-evi)# rd 111.111.111.111:65535
```

Supported releases 10.4.2.0 or later

route-target

Configures the route target (RT) values used in EVPN routes.

Syntax `route-target {auto | value {import | export | both} [asn4]}`

Parameters

value **{import | export | both}**	Configure an RT import or export value, or both values, in the format *2-octet-ASN*:*4-octet-number* or *4-octet-ASN*:*2-octet-number*. • The *2-octet* ASN or number is 1 to 65535. • The *4-octet* ASN or number is 1 to 4294967295.
auto	Configure the RT import and export values to be automatically generated.
asn4	(Optional) Advertises a 4-byte AS number in route target values.

Default Not configured

Command mode EVPN-EVI

Usage information A route target determines how EVPN routes are distributed among EVPN instances. Configure each RT with an import and export value. When the EVPN routes are advertised, the RT export value configured for export is

attached to each route. The receiving VTEP compares a route's export value with the local RT's import value. If the values match, the routes are downloaded and installed on the VTEP.

- For 2-byte autonomous systems, the RT auto-configures as Type 0 from the 2-byte AS and the 3-byte VNI (Type encoded as 0x0002).
- For 4-byte autonomous systems, the RT auto-configures as Type 2 from the 4-byte AS and the 2-byte EVI (Type encoded as 0x0202).

Example

```
OS10(config)# evpn
OS10(config-evpn)# evi 10
OS10(config-evpn-evi)# vni 10000
OS10(config-evpn-evi)# rd 111.111.111.111:65535
OS10(config-evpn-evi)# route-target 1:3 both
```

Supported releases 10.4.2.0 or later

show evpn evi

Displays the configuration settings of EVPN instances.

Syntax `show evpn evi [id]`

Parameters *id* — (Optional) Enter the EVPN instance ID (1 to 65535).

Default Not configured

Command mode EXEC

Usage information Use the `show evpn evi` command to verify EVPN instance status, associated VXLAN virtual networks and the RD and RT values used in BGP EVPN routes in the EVI. The status of integrated routing and bridging (IRB) and the VRF used for EVPN traffic are also displayed.

Example

```
OS10# show evpn evi 101
EVI : 101, State : up
  Bridge-Domain       : Virtual-Network 101, VNI 101
  Route-Distinguisher : 1:95.0.0.4:101(auto)
  Route-Targets       : 0:101:268435556(auto) both
  Inclusive Multicast : 95.0.0.3
  IRB                 : Enabled(VRF: default)
```

Supported releases 10.4.2.0 or later

show evpn mac

Displays BGP EVPN routes for host MAC addresses.

Syntax `show evpn mac {count | mac-address nn.nn.nn.nn | evi id [mac-address nn.nn.nn.nn | count | next-hop ip-address count]}`

Parameters

- `count` — Displays the total number of local and remote host MAC addresses in EVPN instances.
- `mac-address nn.nn.nn.nn` — Displays the BGP EVPN routes for a specific 48-bit host MAC address.
- `evi id` — Displays the host MAC addresses and next hops in a specified EVPN instance (1 to 65535). To filter the output, you can display information on the host MAC address count for an EVPN ID or for a next-hop IP address, and BGP routes for a specified MAC address.

Default Not configured

Command mode EXEC

Usage information Use the `show evpn mac` command to display the BGP routes for host MAC addresses in EVPN instances.

Examples

```
OS10# show evpn mac
Type  -(lcl): Local (rmt): remote

EVI  Mac-Address        Type  Seq-No  Interface/Next-Hop
50   00:00:00:aa:aa:aa  rmt   0       55.1.1.3
```

```
OS10# show evpn mac count

Total MAC Entries :
  Local MAC Address Count  :    2
  Remote MAC Address Count :    5
```

```
OS10# show evpn mac evi 811 count

EVI 811 MAC Entries :
  Local MAC Address Count  :    1
  Remote MAC Address Count :    2
```

```
OS10# show evpn mac evi 811 next-hop 80.80.1.8 count

EVI 811 next-hop 80.80.1.8 MAC Entries :
  Remote MAC Address Count :    2
```

Supported releases 10.4.2.0 or later

show evpn mac-ip

Displays the BGP EVPN Type 2 routes used for host MAC-IP address binding.

Syntax `show evpn mac-ip [count | evi evi [mac-address mac-address] | mac-address mac-address | next-hop ip-address]`

Parameters

- `count` — Displays the total number of MAC addresses in EVPN MAC-IP address binding.
- `evi evi` — Enter an EVPN instance ID (1 to 65535).
- `host ip-address` — Enter the IP address of a host that communicates through EVPN routes.
- `mac-address mac-address` — Enter the MAC address of a host that communicates through EVPN routes in the format *nn:nn:nn:nn:nn*.
- `next-hop ip-address` — Enter the IP address of a next-hop switch.

Default Not configured

Command mode EXEC

Usage information Use the `show evpn mac-ip` command to view the MAC-IP address binding for host communication in VXLAN tenant segments.

Example

```
OS10# show evpn mac-ip

Type  -(lcl): Local (rmt): remote

EVI     Mac-Address        Type   Seq-No  Host-IP     Interface/Next-Hop
101     14:18:77:0c:e5:a3  rmt    0       11.11.11.3      95.0.0.5
101     14:18:77:0c:e5:a3  rmt    0       2001:11::11:3   95.0.0.5
101     14:18:77:25:4e:84  rmt    0       55.55.55.1      95.0.0.3
101     14:18:77:25:6f:84  lcl    0       11.11.11.2
101     14:18:77:25:6f:84  lcl    0       2001:11::11:2
102     14:18:77:0c:e5:a4  rmt    0       12.12.12.3      95.0.0.5
```

```
102     14:18:77:0c:e5:a4  rmt    0       2001:12::12:3   95.0.0.5
102     14:18:77:25:4d:b9  rmt    0       12.12.12.1      95.0.0.3
102     14:18:77:25:6e:b9  lcl    0       12.12.12.2
103     14:18:77:25:4e:84  rmt    0       13.13.13.1      95.0.0.3
103     14:18:77:25:4e:84  rmt    0       2001:13::13:1   95.0.0.3
103     14:18:77:25:6f:84  lcl    0       13.13.13.2
103     14:18:77:25:6f:84  lcl    0       2001:13::13:2
104     14:18:77:25:4d:b9  rmt    0       14.14.14.1      95.0.0.3
104     14:18:77:25:4d:b9  rmt    0       2001:14::14:1   95.0.0.3
104     14:18:77:25:6e:b9  lcl    0       14.14.14.2
104     14:18:77:25:6e:b9  lcl    0       2001:14::14:2
105     14:18:77:25:4d:b9  rmt    0       15.15.15.1      95.0.0.3
105     14:18:77:25:4d:b9  rmt    0       2001:15::15:1   95.0.0.3
105     14:18:77:25:6e:b9  lcl    0       15.15.15.2
105     14:18:77:25:6e:b9  lcl    0       2001:15::15:2
106     14:18:77:25:4e:84  rmt    0       16.16.16.1      95.0.0.3
106     14:18:77:25:4e:84  rmt    0       2001:16::16:1   95.0.0.3
106     14:18:77:25:6f:84  lcl    0       16.16.16.2
106     14:18:77:25:6f:84  lcl    0       2001:16::16:2
```

```
OS10# show evpn mac-ip evi 104

Type  -(lcl): Local (rmt): remote

EVI     Mac-Address        Type  Seq-No  Host-IP      Interface/Next-Hop
104     14:18:77:25:4d:b9  rmt   0       14.14.14.1     95.0.0.3
104     14:18:77:25:4d:b9  rmt   0       2001:14::14:1  95.0.0.3
104     14:18:77:25:6e:b9  lcl   0       14.14.14.2
104     14:18:77:25:6e:b9  lcl   0       2001:14::14:2
```

```
OS10# show evpn mac-ip evi 101 mac-address 14:18:77:0c:e5:a3

Type  -(lcl): Local (rmt): remote

EVI     Mac-Address        Type  Seq-No  Host-IP      Interface/Next-Hop
101     14:18:77:0c:e5:a3  rmt   0       11.11.11.3     95.0.0.5
101     14:18:77:0c:e5:a3  rmt   0       2001:11::11:3  95.0.0.5
```

```
OS10# show evpn mac-ip mac-address 14:18:77:25:4e:84

Type  -(lcl): Local (rmt): remote

EVI     Mac-Address        Type   Seq-No  Host-IP      Interface/Next-Hop
101     14:18:77:25:4e:84  rmt    0       55.55.55.1      95.0.0.3
103     14:18:77:25:4e:84  rmt    0       13.13.13.1      95.0.0.3
103     14:18:77:25:4e:84  rmt    0       2001:13::13:1   95.0.0.3
106     14:18:77:25:4e:84  rmt    0       16.16.16.1      95.0.0.3
106     14:18:77:25:4e:84  rmt    0       2001:16::16:1   95.0.0.3
```

Supported releases 10.4.3.0 or later

show evpn vrf

Displays the VRF instances used to forward EVPN routes in VXLAN overlay networks.

Syntax `show evpn vrf [vrf-name]`

Parameters *vrf-name* — (Optional) Enter the name of a non-default tenant VRF instance.

Default Not configured

Command mode EXEC

Usage information Use the `show evpn vrf` command to verify the tenant VRF instances used in EVPN instances to exchange BGP EVPN routes in VXLANs.

Example

```
show evpn vrf

VXLAN-VNI   EVI     Virtual-Network-Instance  VRF-Name
102         102     102                       blue
103         103     103                       default
104         104     104                       blue
106         106     106                       default
105         105     105                       blue
101         101     101                       default
```

Supported releases 10.4.3.0 or later

show evpn vxlan-vni

Displays the VXLAN overlay network for EVPN instances.

Syntax `show evpn vxlan-vni [vni]`

Parameters *vni* — (Optional) Enter the VXLAN virtual-network ID (1 to 16,777,215).

Default Not configured

Command mode EXEC

Usage information Use the `show evpn vxlan-vni` command to verify the VXLAN virtual network and bridge domain used by an EVPN instance.

Example

```
OS10# show evpn vxlan-vni

VXLAN-VNI   EVI     Bridge-Domain
100         65447   65447
```

Supported releases 10.4.2.0 or later

vni

Associates an EVPN instance with a VXLAN network ID.

Syntax `vni vni`

Parameters

vni Enter the virtual-network ID (1 to 16,777,215).

Default Not configured

Command mode EVPN-EVI

Usage information Use the `vni` command in EVPN-EVI mode to configure an EVPN instance with RD and RT values to an overlay VXLAN virtual network.

Example

```
OS10(config)# evpn
OS10(config-evpn)# evi 10
OS10(config-evpn-evi)# vni 10000
```

Supported releases 10.4.2.0 or later

Example: VXLAN with BGP EVPN

This VXLAN with BGP EVPN example uses a typical Clos leaf-spine topology with VXLAN tunnel endpoints (VTEPs). Individual switch configuration shows how to set up an end-to-end VXLAN. eBGP is used for exchanging IP routes in the IP underlay network, and EVPN routes in the VXLAN overlay network. All spine nodes are in one autonomous system — AS 65001. All leaf nodes are in another autonomous system — AS 65002.

- On VTEPs 1 and 2: Access ports are assigned to the virtual network using a switch-scoped VLAN. EVPN is configured using auto-EVI mode.
- On VTEP 3: Access ports are assigned to the virtual network using a port-scoped VLAN. The EVPN instance is configured using manual configuration mode. The Route Distinguisher and Route Target are configured using auto-EVI mode.
- On VTEP 4: Access ports are assigned to the virtual network using a port-scoped VLAN. EVPN is configured using manual configuration mode, including the EVPN instance, Route Distinguisher, and Route Target.

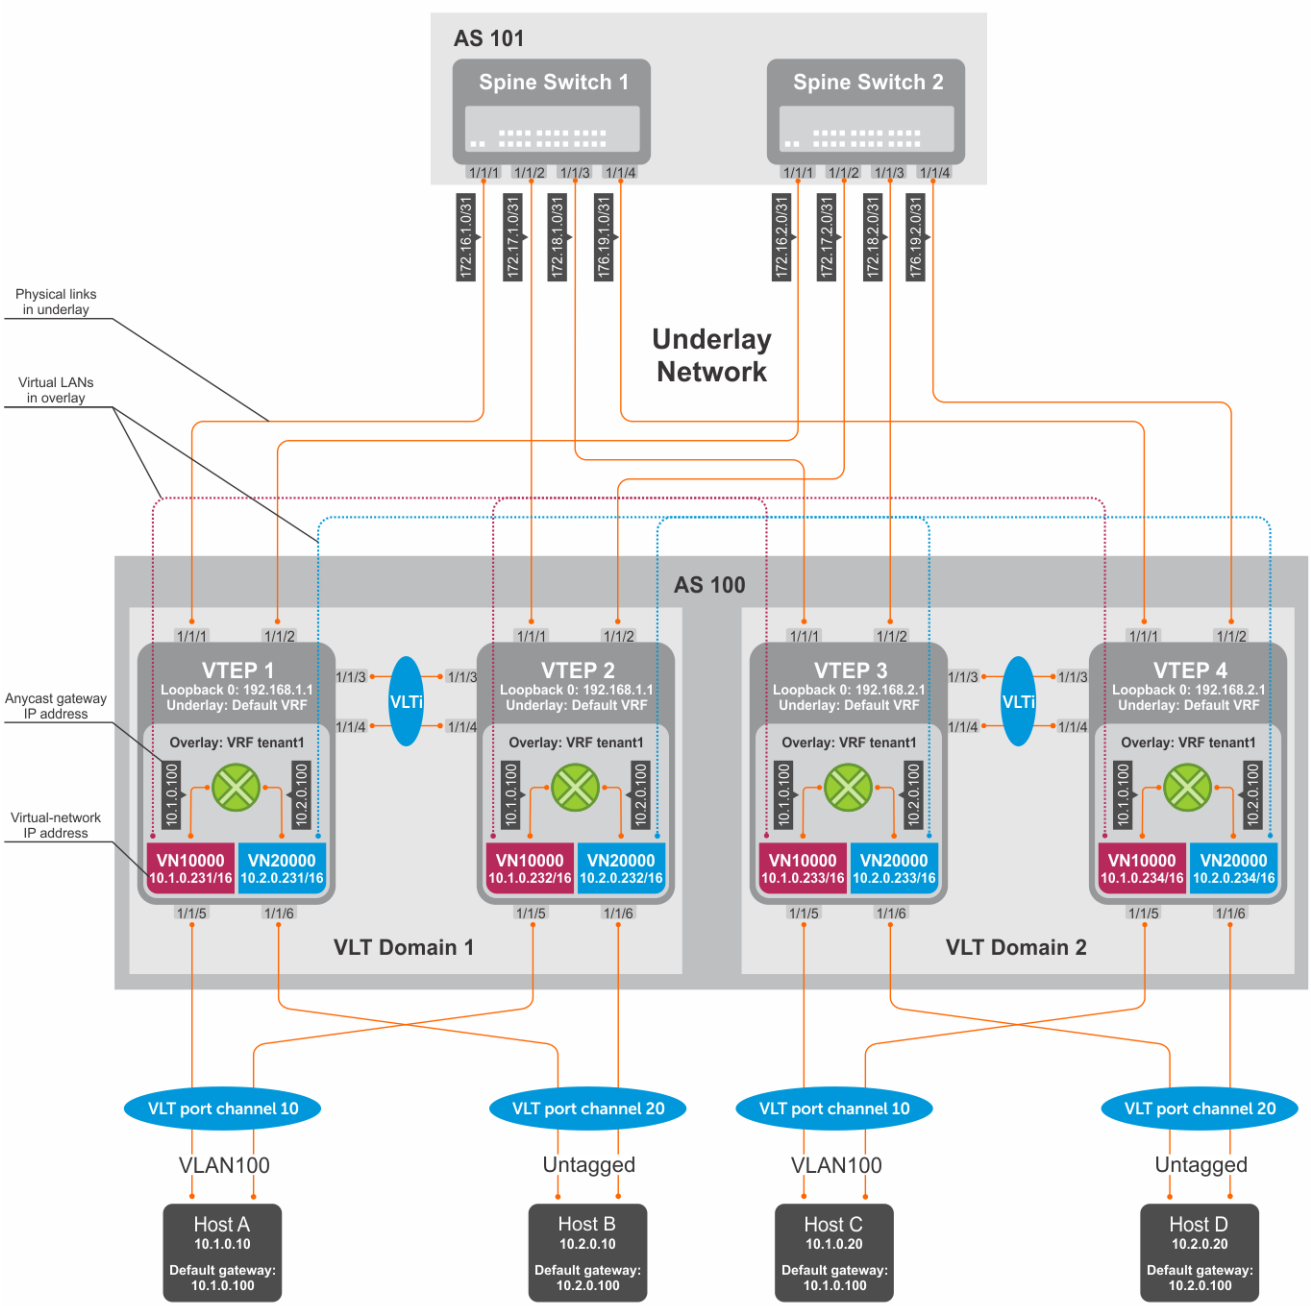


Figure 9. VXLAN BGP EVPN use case

VTEP 1 Leaf Switch

1. Configure a loopback interface for the VXLAN underlay using same IP address as VLT peer

```
OS10(config)# interface loopback0
OS10(conf-if-lo-0)# no shutdown
OS10(conf-if-lo-0)# ip address 192.168.1.1/32
OS10(conf-if-lo-0)# exit
```

2. Configure the loopback interface as the VXLAN source tunnel interface

```
OS10(config)# nve
OS10(config-nve)# source-interface loopback0
OS10(config-nve)# exit
```

3. Configure VXLAN virtual networks

```
OS10(config)# virtual-network 10000
OS10(config-vn)# vxlan-vni 10000
OS10(config-vn-vxlan-vni-10000)# exit
OS10(config-vn)# exit
OS10(config)# virtual-network 20000
OS10(config-vn)# vxlan-vni 20000
OS10(config-vn-vxlan-vni-20000)# exit
OS10(config-vn)# exit
```

4. Assign VLAN member interfaces to virtual networks

Use a switch-scoped VLAN-to-VNI mapping:

```
OS10(config)# interface vlan100
OS10(config-if-vl-100)# virtual-network 10000
OS10(config-if-vl-100)# no shutdown
OS10(config-if-vl-100)# exit
OS10(config)# interface vlan200
OS10(config-if-vl-200)# virtual-network 20000
OS10(config-if-vl-200)# no shutdown
OS10(config-if-vl-200)# exit
```

5. Configure access ports as VLAN members for a switch-scoped VLAN-to-VNI mapping

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# no shutdown
OS10(conf-if-po-10)# switchport mode trunk
OS10(conf-if-po-10)# switchport trunk allowed vlan 100
OS10(conf-if-po-10)# exit

OS10(config)# interface ethernet1/1/5
OS10(conf-if-eth1/1/5)# no shutdown
OS10(conf-if-eth1/1/5)# channel-group 10 mode active
OS10(conf-if-eth1/1/5)# no switchport
OS10(conf-if-eth1/1/5)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# no shutdown
OS10(conf-if-po-20)# switchport mode trunk
OS10(conf-if-po-20)# switchport access vlan 200
OS10(conf-if-po-20)# exit

OS10(config)# interface ethernet1/1/6
OS10(conf-if-eth1/1/6)# no shutdown
OS10(conf-if-eth1/1/6)# channel-group 20 mode active
OS10(conf-if-eth1/1/6)# no switchport
OS10(conf-if-eth1/1/6)# exit
```

6. Configure upstream network-facing ports

```
OS10(config)# interface ethernet1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/1)# ip address 172.16.1.0/31
OS10(conf-if-eth1/1/1)# exit

OS10(config)# interface ethernet1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
```

```
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/2)# ip address 172.16.2.0/31
OS10(conf-if-eth1/1/2)# exit
```

7. Configure eBGP

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# address-family ipv4 unicast
OS10(configure-router-bgpv4-af)# redistribute connected
OS10(configure-router-bgpv4-af)# exit
```

8. Configure eBGP for the IPv4 point-to-point peering

```
OS10(config-router-bgp-100)# neighbor 172.16.1.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-bgp-neighbor-af)# allowas-in 1
OS10(config-router-bgp-neighbor-af)# exit
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit

OS10(config-router-bgp-100)# neighbor 172.16.2.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-bgp-neighbor-af)# allowas-in 1
OS10(config-router-neighbor)# exit
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# exit
```

9. Configure a loopback interface for BGP EVPN peering (different from VLT peer's IP address)

```
OS10(config)# interface loopback1
OS10(conf-if-lo-1)# no shutdown
OS10(conf-if-lo-1)# ip address 172.16.0.1/32
OS10(conf-if-lo-1)# exit
```

10. Configure BGP EVPN peering

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# neighbor 172.201.0.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# ebgp-multihop 4
OS10(config-router-neighbor)# send-community extended
OS10(config-router-neighbor)# update-source loopback1
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-neighbor-af)# no activate
OS10(config-router-neighbor-af)# exit
OS10(config-router-bgp-neighbor)# address-family l2vpn evpn
OS10(config-router-neighbor-af)# activate
OS10(config-router-neighbor-af)# allowas-in 1
OS10(config-router-neighbor-af)# exit
OS10(config-router-neighbor)# exit

OS10(config-router-bgp-100)# neighbor 172.202.0.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# ebgp-multihop 4
OS10(config-router-neighbor)# send-community extended
OS10(config-router-neighbor)# update-source loopback1
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-neighbor-af)# no activate
OS10(config-router-neighbor-af)# exit
OS10(config-router-neighbor)# address-family l2vpn evpn
OS10(config-router-neighbor-af)# activate
OS10(config-router-neighbor-af)# allowas-in 1
OS10(config-router-neighbor-af)# exit
```

```
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# exit
```

11. Configure EVPN

Configure the EVPN instance, Route Distinguisher, and Route Target using auto-EVI mode:

```
OS10(config)# evpn
OS10(config-evpn)# auto-evi
OS10(config-evpn)# exit
```

12. Configure VLT

Configure dedicated L3 underlay path to reach VLT Peer in case of network failure

```
OS10(config)# interface vlan4000
OS10(config-if-vl-4000)# no shutdown
OS10(config-if-vl-4000)# ip address 172.16.250.0/30
OS10(config-if-vl-4000)# ip 1 area 0.0.0.0
OS10(config-if-vl-4000)# exit
```

Configure VLT port channel

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# vlt-port-channel 10
OS10(conf-if-po-10)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# vlt-port-channel 20
OS10(conf-if-po-20)# exit
```

Configure VLTi member links

```
OOS10(config)# interface ethernet1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# exit

OS10(config)# interface ethernet1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# exit
```

Configure VLT domain

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination 10.16.150.1
OS10(conf-vlt-1)# discovery-interface ethernet1/1/3,1/1/4
OS10(conf-vlt-1)# vlt-mac aa:bb:cc:dd:ee:ff
OS10(conf-vlt-1)# exit
```

Configure UFD with uplink VLT ports and downlink network ports

```
OS10(config)# uplink-state-group 1
OS10(conf-uplink-state-group-1)# enable
OS10(conf-uplink-state-group-1)# downstream ethernet1/1/1-1/1/2
OS10(conf-uplink-state-group-1)# upstream port-channel10
OS10(conf-uplink-state-group-1)# upstream port-channel20
OS10(conf-uplink-state-group-1)# exit
```

Configure iBGP IPv4 peering between VLT peers

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# neighbor 172.16.250.1
OS10(config-router-neighbor)# remote-as 100
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# exit
```

13. Configure IP switching in overlay network

Create tenant VRF

```
OS10(config)# ip vrf tenant1
OS10(conf-vrf)# exit
```

Configure anycast gateway MAC address

```
OS10(config)# ip virtual-router mac-address 00:01:01:01:01:01
```

Configure routing on virtual networks

```
OS10(config)# interface virtual-network10000
OS10(conf-if-vn-10000)# ip vrf forwarding tenant1
OS10(conf-if-vn-10000)# ip address 10.1.0.231/16
OS10(conf-if-vn-10000)# ip virtual-router address 10.1.0.100
OS10(conf-if-vn-10000)# no shutdown
OS10(conf-if-vn-10000)# exit

OS10(config)# interface virtual-network20000
OS10(conf-if-vn-20000)# ip vrf forwarding tenant1
OS10(conf-if-vn-20000)# ip address 10.2.0.231/16
OS10(conf-if-vn-10000)# ip virtual-router address 10.2.0.100
OS10(conf-if-vn-20000)# no shutdown
OS10(conf-if-vn-20000)# exit
```

VTEP 2 Leaf Switch

1. Configure a loopback interface for the VXLAN underlay using same IP address as VLT peer

```
OS10(config)# interface loopback0
OS10(conf-if-lo-0)# no shutdown
OS10(conf-if-lo-0)# ip address 192.168.1.1/32
OS10(conf-if-lo-0)# exit
```

2. Configure the loopback interface as the VXLAN source tunnel interface

```
OS10(config)# nve
OS10(config-nve)# source-interface loopback0
OS10(config-nve)# exit
```

3. Configure VXLAN virtual networks

```
OS10(config)# virtual-network 10000
OS10(config-vn)# vxlan-vni 10000
OS10(config-vn-vxlan-vni-10000)# exit
OS10(config-vn)# exit
OS10(config)# virtual-network 20000
OS10(config-vn)# vxlan-vni 20000
OS10(config-vn-vxlan-vni-20000)# exit
OS10(config-vn)# exit
```

4. Assign VLAN member interfaces to virtual networks

Use a switch-scoped VLAN-to-VNI mapping:

```
OS10(config)# interface vlan100
OS10(config-if-vl-100)# virtual-network 10000
OS10(config-if-vl-100)# no shutdown
OS10(config-if-vl-100)# exit
OS10(config)# interface vlan200
OS10(config-if-vl-200)# virtual-network 20000
OS10(config-if-vl-200)# no shutdown
OS10(config-if-vl-200)# exit
```

5. Configure access ports as VLAN members for a switch-scoped VLAN-to-VNI mapping

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# no shutdown
OS10(conf-if-po-10)# switchport mode trunk
OS10(conf-if-po-10)# switchport trunk allowed vlan 100
OS10(conf-if-po-10)# exit

OS10(config)# interface ethernet1/1/5
OS10(conf-if-eth1/1/5)# no shutdown
OS10(conf-if-eth1/1/5)# channel-group 10 mode active
OS10(conf-if-eth1/1/5)# no switchport
OS10(conf-if-eth1/1/5)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# no shutdown
OS10(conf-if-po-20)# switchport mode trunk
OS10(conf-if-po-20)# switchport access vlan 200
OS10(conf-if-po-20)# exit

OS10(config)# interface ethernet1/1/6
OS10(conf-if-eth1/1/6)# no shutdown
OS10(conf-if-eth1/1/6)# channel-group 20 mode active
OS10(conf-if-eth1/1/6)# no switchport
OS10(conf-if-eth1/1/6)# exit
```

6. Configure upstream network-facing ports

```
OS10(config)# interface ethernet1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/1)# ip address 172.17.1.0/31
OS10(conf-if-eth1/1/1)# exit

OS10(config)# interface ethernet1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/2)# ip address 172.17.2.0/31
OS10(conf-if-eth1/1/2)# exit
```

7. Configure eBGP

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# address-family ipv4 unicast
OS10(configure-router-bgpv4-af)# redistribute connected
OS10(configure-router-bgpv4-af)# exit
```

8. Configure eBGP for the IPv4 point-to-point peering

```
OS10(config-router-bgp-100)# neighbor 172.17.1.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-bgp-neighbor-af)# allowas-in 1
OS10(config-router-bgp-neighbor-af)# exit
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit

OS10(config-router-bgp-100)# neighbor 172.17.2.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-bgp-neighbor-af)# allowas-in 1
OS10(config-router-neighbor)# exit
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# exit
```

9. Configure a loopback interface for BGP EVPN peering (different from VLT peer's IP address)

```
OS10(config)# interface loopback1
OS10(conf-if-lo-1)# no shutdown
OS10(conf-if-lo-1)# ip address 172.17.0.1/32
OS10(conf-if-lo-1)# exit
```

10. Configure BGP EVPN peering

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# neighbor 172.201.0.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# ebgp-multihop 4
OS10(config-router-neighbor)# send-community extended
OS10(config-router-neighbor)# update-source loopback1
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-neighbor-af)# no activate
OS10(config-router-neighbor-af)# exit
OS10(config-router-bgp-neighbor)# address-family l2vpn evpn
OS10(config-router-neighbor-af)# activate
OS10(config-router-neighbor-af)# allowas-in 1
OS10(config-router-neighbor-af)# exit
OS10(config-router-neighbor)# exit

OS10(config-router-bgp-100)# neighbor 172.202.0.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# ebgp-multihop 4
OS10(config-router-neighbor)# send-community extended
OS10(config-router-neighbor)# update-source loopback1
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-neighbor-af)# no activate
OS10(config-router-neighbor-af)# exit
OS10(config-router-neighbor)# address-family l2vpn evpn
OS10(config-router-neighbor-af)# activate
OS10(config-router-neighbor-af)# allowas-in 1
OS10(config-router-neighbor-af)# exit
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# exit
```

11. Configure EVPN

Configure the EVPN instance, Route Distinguisher, and Route Target using auto-EVI mode:

```
OS10(config)# evpn
OS10(config-evpn)# auto-evi
OS10(config-evpn)# exit
```

12. Configure VLT

Configure dedicated L3 underlay path to reach VLT Peer in case of network failure

```
OS10(config)# interface vlan4000
OS10(config-if-vl-4000)# no shutdown
OS10(config-if-vl-4000)# ip address 172.16.250.1/30
OS10(config-if-vl-4000)# ip 1 area 0.0.0.0
OS10(config-if-vl-4000)# exit
```

Configure VLT port channel

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# vlt-port-channel 10
OS10(conf-if-po-10)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# vlt-port-channel 20
OS10(conf-if-po-20)# exit
```

Configure VLTi member links

```
OOS10(config)# interface ethernet1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# exit

OS10(config)# interface ethernet1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# exit
```

Configure VLT domain

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination 10.16.150.2
OS10(conf-vlt-1)# discovery-interface ethernet1/1/3,1/1/4
OS10(conf-vlt-1)# vlt-mac aa:bb:cc:dd:ee:ff
OS10(conf-vlt-1)# exit
```

Configure UFD with uplink VLT ports and downlink network ports

```
OS10(config)# uplink-state-group 1
OS10(conf-uplink-state-group-1)# enable
OS10(conf-uplink-state-group-1)# downstream ethernet1/1/1-1/1/2
OS10(conf-uplink-state-group-1)# upstream port-channel10
OS10(conf-uplink-state-group-1)# upstream port-channel20
OS10(conf-uplink-state-group-1)# exit
```

Configure iBGP IPv4 peering between VLT peers

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# neighbor 172.16.250.0
OS10(config-router-neighbor)# remote-as 100
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# exit
```

13. Configure IP switching in overlay network

Create tenant VRF

```
OS10(config)# ip vrf tenant1
OS10(conf-vrf)# exit
```

Configure anycast gateway MAC address

```
OS10(config)# ip virtual-router mac-address 00:01:01:01:01:01
```

Configure routing on virtual networks

```
OS10(config)# interface virtual-network10000
OS10(conf-if-vn-10000)# ip vrf forwarding tenant1
OS10(conf-if-vn-10000)# ip address 10.1.0.232/16
OS10(conf-if-vn-10000)# ip virtual-router address 10.1.0.100
OS10(conf-if-vn-10000)# no shutdown
OS10(conf-if-vn-10000)# exit

OS10(config)# interface virtual-network20000
OS10(conf-if-vn-20000)# ip vrf forwarding tenant1
OS10(conf-if-vn-20000)# ip address 10.2.0.232/16
OS10(conf-if-vn-10000)# ip virtual-router address 10.2.0.100
OS10(conf-if-vn-20000)# no shutdown
OS10(conf-if-vn-20000)# exit
```

VTEP 3 Leaf Switch

1. Configure a loopback interface for the VXLAN underlay using same IP address as VLT peer

```
OS10(config)# interface loopback0
OS10(conf-if-lo-0)# no shutdown
OS10(conf-if-lo-0)# ip address 192.168.2.1/32
OS10(conf-if-lo-0)# exit
```

2. Configure the loopback interface as the VXLAN source tunnel interface

```
OS10(config)# nve
OS10(config-nve)# source-interface loopback0
OS10(config-nve)# exit
```

3. Configure VXLAN virtual networks

```
OS10(config)# virtual-network 10000
OS10(config-vn)# vxlan-vni 10000
OS10(config-vn-vxlan-vni-10000)# exit
OS10(config-vn)# exit

OS10(config)# virtual-network 20000
OS10(config-vn)# vxlan-vni 20000
OS10(config-vn-vxlan-vni-20000)# exit
OS10(config-vn)# exit
```

4. Configure unused VLAN ID for untagged membership

```
OS10(config)# virtual-network untagged-vlan 4000
```

5. Configure access ports as VLAN members for a port-scoped VLAN-to-VNI mapping

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# no shutdown
OS10(conf-if-po-10)# switchport mode trunk
OS10(conf-if-po-10)# exit

OS10(config)# interface ethernet1/1/5
OS10(conf-if-eth1/1/5)# no shutdown
OS10(conf-if-eth1/1/5)# channel-group 10 mode active
OS10(conf-if-eth1/1/5)# no switchport
OS10(conf-if-eth1/1/5)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# no shutdown
OS10(conf-if-po-20)# switchport mode trunk
OS10(conf-if-po-20)# exit

OS10(config)# interface ethernet1/1/6
OS10(conf-if-eth1/1/6)# no shutdown
OS10(conf-if-eth1/1/6)# channel-group 20 mode active
OS10(conf-if-eth1/1/6)# no switchport
OS10(conf-if-eth1/1/6)# exit
```

6. Add access ports to virtual networks

```
OS10(config)# virtual-network 10000
OS10(config-vn)# member-interface port-channel 10 vlan-tag 100
OS10(config-vn)# exit

OS10(config)# virtual-network 20000
OS10(config-vn)# member-interface ethernet 1/1/6 untagged
OS10(config-vn)# exit
```

7. Configure upstream network-facing ports

```
OS10(config)# interface ethernet1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/1)# ip address 172.18.1.0/31
OS10(conf-if-eth1/1/1)# exit

OS10(config)# interface ethernet1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/2)# ip address 172.18.2.0/31
OS10(conf-if-eth1/1/2)# exit
```

8. Configure eBGP

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# address-family ipv4 unicast
OS10(configure-router-bgpv4-af)# redistribute connected
OS10(configure-router-bgpv4-af)# exit
```

9. Configure eBGP for the IPv4 point-to-point peering

```
OS10(config-router-bgp-100)# neighbor 172.18.1.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-bgp-neighbor-af)# allowas-in 1
OS10(config-router-bgp-neighbor-af)# exit
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit

OS10(config-router-bgp-100)# neighbor 172.18.2.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-bgp-neighbor-af)# allowas-in 1
OS10(config-router-neighbor)# exit
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# exit
```

10. Configure a loopback interface for BGP EVPN peering (different from VLT peer's IP address)

```
OS10(config)# interface loopback1
OS10(conf-if-lo-1)# no shutdown
OS10(conf-if-lo-1)# ip address 172.18.0.1/32
OS10(conf-if-lo-1)# exit
```

11. Configure BGP EVPN peering

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# neighbor 172.201.0.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# ebgp-multihop 4
OS10(config-router-neighbor)# send-community extended
OS10(config-router-neighbor)# update-source loopback1
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-neighbor-af)# no activate
OS10(config-router-neighbor-af)# exit
OS10(config-router-bgp-neighbor)# address-family l2vpn evpn
OS10(config-router-neighbor-af)# activate
OS10(config-router-neighbor-af)# allowas-in 1
OS10(config-router-neighbor-af)# exit
OS10(config-router-neighbor)# exit

OS10(config-router-bgp-100)# neighbor 172.202.0.1
```

```
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# ebgp-multihop 4
OS10(config-router-neighbor)# send-community extended
OS10(config-router-neighbor)# update-source loopback1
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-neighbor-af)# no activate
OS10(config-router-neighbor-af)# exit
OS10(config-router-neighbor)# address-family l2vpn evpn
OS10(config-router-neighbor-af)# activate
OS10(config-router-neighbor-af)# allowas-in 1
OS10(config-router-neighbor-af)# exit
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# exit
```

12. Configure EVPN

Configure the EVPN instance, Route Distinguisher, and Route Target using manual configuration mode:

```
OS10(config)# evpn
OS10(config-evpn)# evi 1000
OS10(config-evpn-evi)# vni 10000
OS10(config-evpn-evi)# rd auto
OS10(config-evpn-evi)# route-target both auto
OS10(config-evpn-evi)# exit

OS10(config-evpn)# evi 2000
OS10(config-evpn-evi)# vni 20000
OS10(config-evpn-evi)# rd auto
OS10(config-evpn-evi)# route-target both auto
OS10(config-evpn-evi)# exit
OS10(config-evpn)# exit
```

13. Configure VLT

Configure VLTi VLAN for virtual network

```
OS10(config)# virtual-network 10000
OS10(config-vn)# vlti-vlan 100
OS10(config-vn)# exit
```

Configure dedicated L3 underlay path to reach VLT Peer in case of network failure

```
OS10(config)# interface vlan4000
OS10(config-if-vl-4000)# no shutdown
OS10(config-if-vl-4000)# ip address 172.16.250.10/30
OS10(config-if-vl-4000)# exit
```

Configure VLT port channels

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# vlt-port-channel 10
OS10(conf-if-po-10)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# vlt-port-channel 20
OS10(conf-if-po-20)# exit
```

Configure VLTi member links

```
OOS10(config)# interface ethernet1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# exit

OS10(config)# interface ethernet1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
```

```
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# exit
```

Configure VLT domain

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination 10.16.150.3
OS10(conf-vlt-1)# discovery-interface ethernet1/1/3,1/1/4
OS10(conf-vlt-1)# vlt-mac aa:bb:cc:dd:ee:ff
OS10(conf-vlt-1)# exit
```

Configure UFD with uplink VLT ports and downlink network ports

```
OS10(config)# uplink-state-group 1
OS10(conf-uplink-state-group-1)# enable
OS10(conf-uplink-state-group-1)# downstream ethernet1/1/1-1/1/2
OS10(conf-uplink-state-group-1)# upstream port-channel10
OS10(conf-uplink-state-group-1)# upstream port-channel20
OS10(conf-uplink-state-group-1)# exit
```

Configure iBGP IPv4 peering between VLT peers

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# neighbor 172.16.250.11
OS10(config-router-neighbor)# remote-as 100
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# exit
```

14. Configure IP routing in overlay network

Create tenant VRF

```
OS10(config)# ip vrf tenant1
OS10(conf-vrf)# exit
```

Configure anycast gateway MAC address

```
OS10(config)# ip virtual-router mac-address 00:01:01:01:01:01
```

Configure routing on virtual networks

```
OS10(config)# interface virtual-network10000
OS10(conf-if-vn-10000)# ip vrf forwarding tenant1
OS10(conf-if-vn-10000)# ip address 10.1.0.233/16
OS10(conf-if-vn-10000)# ip virtual-router address 10.1.0.100
OS10(conf-if-vn-10000)# no shutdown
OS10(conf-if-vn-10000)# exit

OS10(config)# interface virtual-network20000
OS10(conf-if-vn-20000)# ip vrf forwarding tenant1
OS10(conf-if-vn-20000)# ip address 10.2.0.233/16
OS10(conf-if-vn-10000)# ip virtual-router address 10.2.0.100
OS10(conf-if-vn-20000)# no shutdown
OS10(conf-if-vn-20000)# exit
```

VTEP 4 Leaf Switch

1. Configure a loopback interface for the VXLAN underlay using same IP address as VLT peer

```
OS10(config)# interface loopback0
OS10(conf-if-lo-0)# no shutdown
OS10(conf-if-lo-0)# ip address 192.168.2.1/32
OS10(conf-if-lo-0)# exit
```

2. Configure the loopback interface as the VXLAN source tunnel interface

```
OS10(config)# nve
OS10(config-nve)# source-interface loopback0
OS10(config-nve)# exit
```

3. Configure VXLAN virtual networks

```
OS10(config)# virtual-network 10000
OS10(config-vn)# vxlan-vni 10000
OS10(config-vn-vxlan-vni-10000)# exit
OS10(config-vn)# exit

OS10(config)# virtual-network 20000
OS10(config-vn)# vxlan-vni 20000
OS10(config-vn-vxlan-vni-20000)# exit
OS10(config-vn)# exit
```

4. Configure unused VLAN ID for untagged membership

```
OS10(config)# virtual-network untagged-vlan 4000
```

5. Configure access ports as VLAN members for a port-scoped VLAN-to-VNI mapping

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# no shutdown
OS10(conf-if-po-10)# switchport mode trunk
OS10(conf-if-po-10)# exit

OS10(config)# interface ethernet1/1/5
OS10(conf-if-eth1/1/5)# no shutdown
OS10(conf-if-eth1/1/5)# channel-group 10 mode active
OS10(conf-if-eth1/1/5)# no switchport
OS10(conf-if-eth1/1/5)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# no shutdown
OS10(conf-if-po-20)# switchport mode trunk
OS10(conf-if-po-20)# exit

OS10(config)# interface ethernet1/1/6
OS10(conf-if-eth1/1/6)# no shutdown
OS10(conf-if-eth1/1/6)# channel-group 20 mode active
OS10(conf-if-eth1/1/6)# no switchport
OS10(conf-if-eth1/1/6)# exit
```

6. Add access ports to virtual networks

```
OS10(config)# virtual-network 10000
OS10(config-vn)# member-interface port-channel 10 vlan-tag 100
OS10(config-vn)# exit

OS10(config)# virtual-network 20000
OS10(config-vn)# member-interface ethernet 1/1/6 untagged
OS10(config-vn)# exit
```

7. Configure upstream network-facing ports

```
OS10(config)# interface ethernet1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
OS10(conf-if-eth1/1/1)# ip address 172.19.1.0/31
OS10(conf-if-eth1/1/1)# exit

OS10(config)# interface ethernet1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/1)# mtu 1650
```

```
OS10(conf-if-eth1/1/2)# ip address 172.19.2.0/31
OS10(conf-if-eth1/1/2)# exit
```

8. Configure eBGP

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# address-family ipv4 unicast
OS10(configure-router-bgpv4-af)# redistribute connected
OS10(configure-router-bgpv4-af)# exit
```

9. Configure eBGP for the IPv4 point-to-point peering

```
OS10(config-router-bgp-100)# neighbor 172.19.1.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-bgp-neighbor-af)# allowas-in 1
OS10(config-router-bgp-neighbor-af)# exit
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit

OS10(config-router-bgp-100)# neighbor 172.19.2.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-bgp-neighbor-af)# allowas-in 1
OS10(config-router-neighbor)# exit
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# exit
```

10. Configure a loopback interface for BGP EVPN peering (different from VLT peer's IP address)

```
OS10(config)# interface loopback1
OS10(conf-if-lo-1)# no shutdown
OS10(conf-if-lo-1)# ip address 172.19.0.1/32
OS10(conf-if-lo-1)# exit
```

11. Configure BGP EVPN peering

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# neighbor 172.201.0.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# ebgp-multihop 4
OS10(config-router-neighbor)# send-community extended
OS10(config-router-neighbor)# update-source loopback1
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-neighbor-af)# no activate
OS10(config-router-neighbor-af)# exit
OS10(config-router-bgp-neighbor)# address-family l2vpn evpn
OS10(config-router-neighbor-af)# activate
OS10(config-router-neighbor-af)# allowas-in 1
OS10(config-router-neighbor-af)# exit
OS10(config-router-neighbor)# exit

OS10(config-router-bgp-100)# neighbor 172.202.0.1
OS10(config-router-neighbor)# remote-as 101
OS10(config-router-neighbor)# ebgp-multihop 4
OS10(config-router-neighbor)# send-community extended
OS10(config-router-neighbor)# update-source loopback1
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# address-family ipv4 unicast
OS10(config-router-neighbor-af)# no activate
OS10(config-router-neighbor-af)# exit
OS10(config-router-neighbor)# address-family l2vpn evpn
OS10(config-router-neighbor-af)# activate
OS10(config-router-neighbor-af)# allowas-in 1
OS10(config-router-neighbor-af)# exit
```

```
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# exit
```

12. Configure EVPN

Configure the EVPN instance, Route Distinguisher, and Route Target using manual configuration mode:

```
OS10(config)# evpn
OS10(config-evpn)# evi 1000
OS10(config-evpn-evi)# vni 10000
OS10(config-evpn-evi)# rd auto
OS10(config-evpn-evi)# route-target both auto
OS10(config-evpn-evi)# exit

OS10(config-evpn)# evi 2000
OS10(config-evpn-evi)# vni 20000
OS10(config-evpn-evi)# rd auto
OS10(config-evpn-evi)# route-target both auto
OS10(config-evpn-evi)# exit
OS10(config-evpn)# exit
```

13. Configure VLT

Configure VLTi VLAN for virtual network

```
OS10(config)# virtual-network 10000
OS10(config-vn)# vlti-vlan 100
OS10(config-vn)# exit
```

Configure dedicated L3 underlay path to reach VLT Peer in case of network failure

```
OS10(config)# interface vlan4000
OS10(config-if-vl-4000)# no shutdown
OS10(config-if-vl-4000)# ip address 172.16.250.11/30
OS10(config-if-vl-4000)# exit
```

Configure VLT port channels

```
OS10(config)# interface port-channel10
OS10(conf-if-po-10)# vlt-port-channel 10
OS10(conf-if-po-10)# exit

OS10(config)# interface port-channel20
OS10(conf-if-po-20)# vlt-port-channel 20
OS10(conf-if-po-20)# exit
```

Configure VLTi member links

```
OOS10(config)# interface ethernet1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# exit

OS10(config)# interface ethernet1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# exit
```

Configure VLT domain

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination 10.16.150.4
OS10(conf-vlt-1)# discovery-interface ethernet1/1/3,1/1/4
OS10(conf-vlt-1)# vlt-mac aa:bb:cc:dd:ee:ff
OS10(conf-vlt-1)# exit
```

Configure UFD with uplink VLT ports and downlink network ports

```
OS10(config)# uplink-state-group 1
OS10(conf-uplink-state-group-1)# enable
OS10(conf-uplink-state-group-1)# downstream ethernet1/1/1-1/1/2
OS10(conf-uplink-state-group-1)# upstream port-channel10
OS10(conf-uplink-state-group-1)# upstream port-channel20
OS10(conf-uplink-state-group-1)# exit
```

Configure iBGP IPv4 peering between VLT peers

```
OS10(config)# router bgp 100
OS10(config-router-bgp-100)# neighbor 172.16.250.10
OS10(config-router-neighbor)# remote-as 100
OS10(config-router-neighbor)# no shutdown
OS10(config-router-neighbor)# exit
OS10(config-router-bgp-100)# exit
```

14. Configure IP routing in overlay network

Create tenant VRF

```
OS10(config)# ip vrf tenant1
OS10(conf-vrf)# exit
```

Configure anycast gateway MAC address

```
OS10(config)# ip virtual-router mac-address 00:01:01:01:01:01
```

Configure routing on virtual networks

```
OS10(config)# interface virtual-network10000
OS10(conf-if-vn-10000)# ip vrf forwarding tenant1
OS10(conf-if-vn-10000)# ip address 10.1.0.234/16
OS10(conf-if-vn-10000)# ip virtual-router address 10.1.0.100
OS10(conf-if-vn-10000)# no shutdown
OS10(conf-if-vn-10000)# exit

OS10(config)# interface virtual-network20000
OS10(conf-if-vn-20000)# ip vrf forwarding tenant1
OS10(conf-if-vn-20000)# ip address 10.2.0.234/16
OS10(conf-if-vn-10000)# ip virtual-router address 10.2.0.100
OS10(conf-if-vn-20000)# no shutdown
OS10(conf-if-vn-20000)# exit
```

Spine Switch 1

1. Configure downstream ports on underlay links to leaf switches

```
OS10(config)# interface ethernet1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# ip address 172.16.1.1/31
OS10(conf-if-eth1/1/1)# exit
OS10(config)# interface ethernet1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/2)# ip address 172.17.1.1/31
OS10(conf-if-eth1/1/2)# exit
OS10(config)# interface ethernet1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# ip address 172.18.1.1/31
OS10(conf-if-eth1/1/3)# exit
OS10(config)# interface ethernet1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
```

```
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# ip address 172.19.1.1/31
OS10(conf-if-eth1/1/4)# exit
```

2. Configure eBGP

```
OS10(config)# router bgp 101
OS10(config-router-bgp-101)# address-family ipv4 unicast
OS10(config-router-bgp-101)# redistribute connected
```

3. Configure eBGP IPv4 peer sessions on P2P links

```
OS10(conf-router-bgp-101)# neighbor 172.16.1.0
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# exit

OS10(conf-router-bgp-101)# neighbor 172.17.1.0
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# exit

OS10(conf-router-bgp-101)# neighbor 172.18.1.0
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# exit

OS10(conf-router-bgp-101)# neighbor 172.19.1.0
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# exit
OS10(conf-router-bgp-101)# exit
```

4. Configure loopback interface for BGP EVPN peering

```
OS10(config)# interface loopback1
OS10(conf-if-lo-1)# no shutdown
OS10(conf-if-lo-1)# ip address 172.202.0.1/32
OS10(conf-if-lo-1)# exit
```

5. Configure BGP EVPN peer sessions

```
OS10(config)# router bgp 101
OS10(conf-router-bgp-101)# neighbor 172.16.0.1
OS10(conf-router-neighbor)# ebgp-multihop 4
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# send-community extended
OS10(conf-router-neighbor)# update-source loopback1
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no activate
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# address-family l2vpn evpn
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# activate
OS10(conf-router-neighbor-af)# exit
```

```
OS10(conf-router-bgp-101)# neighbor 172.17.0.1
OS10(conf-router-neighbor)# ebgp-multihop 4
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# send-community extended
OS10(conf-router-neighbor)# update-source loopback1
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no activate
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# address-family l2vpn evpn
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# activate
OS10(conf-router-neighbor-af)# exit

OS10(conf-router-bgp-101)# neighbor 172.18.0.1
OS10(conf-router-neighbor)# ebgp-multihop 4
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# send-community extended
OS10(conf-router-neighbor)# update-source loopback1
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no activate
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# address-family l2vpn evpn
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# activate
OS10(conf-router-neighbor-af)# exit

OS10(conf-router-bgp-101)# neighbor 172.19.0.1
OS10(conf-router-neighbor)# ebgp-multihop 4
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# send-community extended
OS10(conf-router-neighbor)# update-source loopback1
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no activate
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# address-family l2vpn evpn
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# activate
OS10(conf-router-neighbor-af)# exit
```

Spine Switch 2

1. Configure downstream ports on underlay links to leaf switches

```
OS10(config)# interface ethernet1/1/1
OS10(conf-if-eth1/1/1)# no shutdown
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# ip address 172.16.2.1/31
OS10(conf-if-eth1/1/1)# exit
OS10(config)# interface ethernet1/1/2
OS10(conf-if-eth1/1/2)# no shutdown
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/2)# ip address 172.17.2.1/31
OS10(conf-if-eth1/1/2)# exit
OS10(config)# interface ethernet1/1/3
OS10(conf-if-eth1/1/3)# no shutdown
OS10(conf-if-eth1/1/3)# no switchport
OS10(conf-if-eth1/1/3)# ip address 172.18.2.1/31
OS10(conf-if-eth1/1/3)# exit
OS10(config)# interface ethernet1/1/4
OS10(conf-if-eth1/1/4)# no shutdown
OS10(conf-if-eth1/1/4)# no switchport
OS10(conf-if-eth1/1/4)# ip address 172.19.2.1/31
OS10(conf-if-eth1/1/4)# exit
```

2. Configure eBGP

```
OS10(config)# router bgp 101
OS10(config-router-bgp-101)# address-family ipv4 unicast
OS10(config-router-bgp-101)# redistribute connected
```

3. Configure eBGP IPv4 peer sessions on P2P links

```
OS10(conf-router-bgp-101)# neighbor 172.16.1.0
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# exit

OS10(conf-router-bgp-101)# neighbor 172.17.1.0
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# exit

OS10(conf-router-bgp-101)# neighbor 172.18.1.0
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# exit

OS10(conf-router-bgp-101)# neighbor 172.19.1.0
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# exit
OS10(conf-router-bgp-101)# exit
```

4. Configure loopback interface for BGP EVPN peering

```
OS10(config)# interface loopback1
OS10(conf-if-lo-1)# no shutdown
OS10(conf-if-lo-1)# ip address 172.202.0.1/32
OS10(conf-if-lo-1)# exit
```

5. Configure BGP EVPN peer sessions

```
OS10(config)# router bgp 101
OS10(conf-router-bgp-101)# neighbor 172.16.0.1
OS10(conf-router-neighbor)# ebgp-multihop 4
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# send-community extended
OS10(conf-router-neighbor)# update-source loopback1
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no activate
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# address-family l2vpn evpn
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# activate
OS10(conf-router-neighbor-af)# exit

OS10(conf-router-bgp-101)# neighbor 172.17.0.1
OS10(conf-router-neighbor)# ebgp-multihop 4
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# send-community extended
```

```
OS10(conf-router-neighbor)# update-source loopback1
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no activate
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# address-family l2vpn evpn
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# activate
OS10(conf-router-neighbor-af)# exit

OS10(conf-router-bgp-101)# neighbor 172.18.0.1
OS10(conf-router-neighbor)# ebgp-multihop 4
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# send-community extended
OS10(conf-router-neighbor)# update-source loopback1
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no activate
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# address-family l2vpn evpn
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# activate
OS10(conf-router-neighbor-af)# exit

OS10(conf-router-bgp-101)# neighbor 172.19.0.1
OS10(conf-router-neighbor)# ebgp-multihop 4
OS10(conf-router-neighbor)# remote-as 100
OS10(conf-router-neighbor)# send-community extended
OS10(conf-router-neighbor)# update-source loopback1
OS10(conf-router-neighbor)# no shutdown
OS10(conf-router-neighbor)# address-family ipv4 unicast
OS10(conf-router-neighbor-af)# no activate
OS10(conf-router-neighbor-af)# exit
OS10(conf-router-neighbor)# address-family l2vpn evpn
OS10(conf-router-neighbor-af)# no sender-side-loop-detection
OS10(conf-router-neighbor-af)# activate
OS10(conf-router-neighbor-af)# exit
```

Controller-provisioned VXLAN

OS10 supports controller-based virtual extensible LAN (VXLAN) implementation using VMware NSX controller, which handles configurations and control plane operations. Controller-based VXLAN deployments are widely deployed solutions for Layer 2 (L2) extension within and between data centers. They are becoming a requirement for cross-data center (DC) connectivity for disaster recovery, storage replication, and general LAN stretching over long distances for virtual appliances and workload mobility. Unlike controller-based VLXAN implementations, in a static VXLAN, an OS10 switch operates as a VXLAN tunnel endpoint (VTEP) and also acts as a VXLAN gateway. The switch performs encapsulation or de-capsulation of L2 frames to be sent to the underlay network; a Layer 3 (L3) network. The VTEP classifies the frames based on manually configured parameters. For more information on static VXLAN configuration, see Example: VXLAN with static VTEP.

NOTE: Controller-based VXLAN is not supported on the S5148F–ON and S3048–ON switches.

Configuration notes

Consider the following when configuring VXLAN in OS10:

- The network virtualization edge (NVE) source interface can only be a Loopback interface, which is a part of the default VRF.
- In an NVE setup, configuring an IPv6 address as the source IP address is not supported.
- OS10 only supports NSX as an OVSDB controller for automatic VXLAN provisioning.
- NSX and other controllers, such as OpenFlow cannot coexist.
- If you reload OS10, you must push the configuration from the controller to an OS10 node.

- The NSX controller supports only service node (SN) replication.
- Implementing combination of VXLAN using the NSX controller, static, and BGP EVPN, is not supported.
- Updating VXLAN access port statistics to the NSX controller is not supported.
- The L2 bridge domain and legacy nonvirtual network LANs can coexist in the same system only when the IDs of VLAN and port-scoped VLAN are different.
- The OS10 switch configured as a VTEP gateway operates with other third-party VTEP gateways to send and receive VXLAN encapsulated traffic.

Standard Compliance

OS10 complies with the RFC5880 for Bidirectional Forwarding Detection (BFD).

VXLAN components

The following table lists the components and the products that are typically deployed in the controller-based VXLAN environment:

Table 23. VXLAN components and products

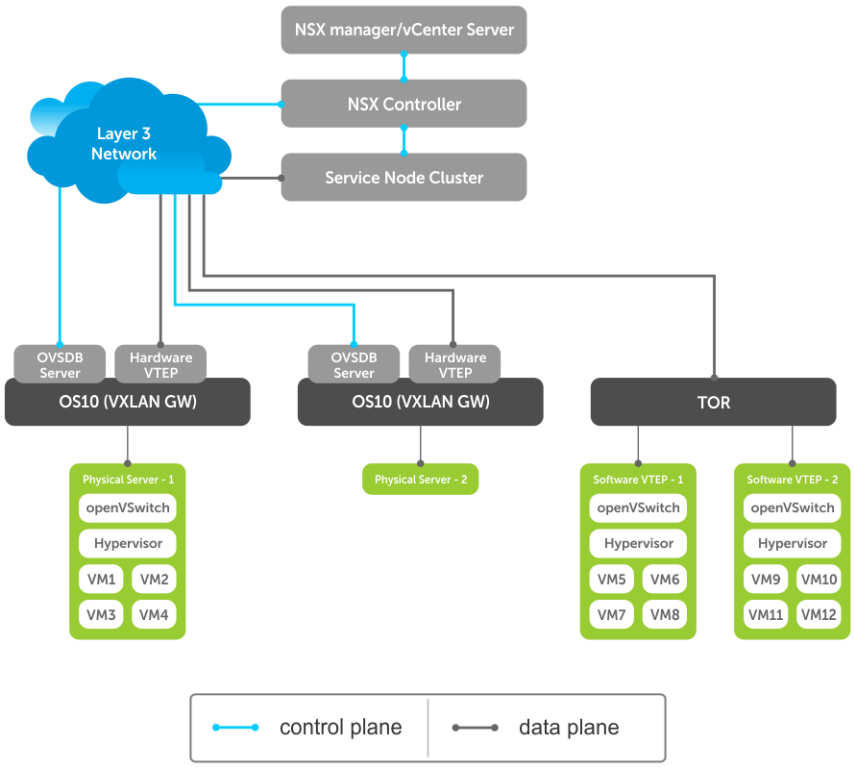
Components	Products
NVP controller	NSX controller
VTEP	OS10 switch
Overlay protocol	VXLAN
Media access control (MAC) learning protocol	Open vSwitch Database (OVSDB)
VXLAN Hypervisor	VMware vSphere
Service Node	VMware ESXi

NVO architecture

- The controller auto-provisions the overlay network. The underlay network is manually configured.
- The controller provisions each virtual network, VTEP, and VNI.
- The controller provisions access port membership in a virtual network.
- The existing routing protocols provision learn the underlay reachability to the VTEP peers.
- The controller learns the remote host MAC addresses. Based on the remote MAC addresses, the controller establishes the unicast tunnel.
- Local MAC addresses learned at the VXLAN access port export to the NSX controller. On aging, the MAC addresses delete from the controller.

OS10 switch and NSX controller integration

The following shows the integration of physical and virtual components in controller-based VXLAN environment:



To communicate server traffic through the extended L2 forwarding domain using VXLAN tunnels, OS10 switch that is deployed as a VXLAN gateway acts as a hardware VTEP. The NSX controller handles control plane operations of the VXLAN. The OS10 VTEP runs an OVSDB server, which learns the MAC address of the connected physical devices and stores it in the hardware VTEP schema. This schema stores the configuration and control information in various OVSDB tables. Based on the schema, the database initializes the OVSDB server.

OS10 supports only the NSX controller in a VXLAN environment. The OVSDB server on the VTEP communicates with the NSX controller through the OVSDB protocol. The VTEP devices and the controller use the schema to exchange information. For example, the VTEP captures the information such as port changes, MAC addresses of the physical network, stores in the schema, and shares with the controller.

In the controller-based VXLAN environment, the NSX controller can perform VTEP operations.

The NSX controller communicates with the OVSDB server in the OS10 VTEP using the OVSDB management protocol over Secure Sockets Layer (SSL) Transmission Control Protocol (TCP) connection. Establishing the communication between the controller and VTEP involves generating the SSL certificate at a VTEP and copying the certificate to the NSX controller. After SSL authentication, a secure connection over SSL establishes between the NSX and the VTEP, and the configuration from NSX manager processes in the VTEP.

Configure the NSX controller

To configure the NSX controller, follow these steps on each OS10 VTEP:

1 Create a Loopback interface and assign an IPv4 address.
2 Create a NVE instance and configure the NVE source interface.
3 Enable an NSX controller.
4 Assign interfaces to expose to the controller.
5 (Optional) Implement Bidirectional Forwarding Detection (BFD) for NSX in the VXLAN network.

For more information about creating a Loopback interface and configuring an NVE source interface, see Configure source IP address on VTEP.

Configure a connection to an OVSDB controller

In OS10 VTEP, the controller configuration command initializes a connection to an OVSDB-based controller. Currently, the only supported OVSDB controller is NSX. The OVSDB controller command also starts the OVSDB server service in the specific OS10 VTEP. OS10 supports only one controller connection at a time.
To configure an OVSDB controller on the OS10 VTEP, follow these steps:

1 Enable VXLAN in CONFIGURATION mode.
```
OS10(config)# nve
```
2 Enable the controller in CONFIGURATION-NVE mode.
```
OS10(config-nve)# controller ovsdb
```
3 Specify the IP address, port of the OVSDB controller, and SSL as a secure connection protocol between the OS10 VTEP and the controller in CONFIGURATION-NVE-OVSDB mode.
```
OS10(config-nve-ovsdb)# ip ip-address port port-number ssl
```
The range is from 0 to 65535. Configure the port-number as `6640` and the connection type as `SSL`.

4 (Optional) Specify the time interval, in milliseconds, that the switch waits between the connection attempts to the controller.
```
OS10(config-nve-ovsdb)# max-backoff interval
```
The range is from 1000 to 180000 milliseconds. The default is 8000 milliseconds.

Sample controller configuration in an OS10 VTEP:

```
OS10# configure terminal
OS10(config)# nve
OS10(config-nve)# controller ovsdb
OS10(config-nve-ovsdb)# ip 10.11.66.110 port 6640 ssl
```

Assign an interface

For the VXLAN setup, explicitly assign interfaces to be managed by an OVSDB controller.
Consider the following before you assign the interface:

- The interface must be in Switchport Trunk mode.
- The interface must not be tagged or untagged to any VLAN
- The interface must not be a member of a port-channel

If these conditions are not met when assigning the interface for controller mode, the system return specific error messages.

When the interface is in controller mode, you cannot:

- remove the interface from Switchport Trunk mode
- tag or untag the interface to any VLAN
- remove the interface from the controller configuration if the interface has active port-scoped VLAN (Port,VLAN) pairs configured by the controller

To assign an interface to Controller mode:

1 Configure an interface from CONFIGURATION mode.
```
OS10(config)# interface ethernet 1/1/1
```
2 Configure L2 trunking in INTERFACE mode.

```
OS10(config-if-eth1/1/1)# switchport mode trunk
```

3 Configure the access VLAN assigned to a L2 trunk port in the INTERFACE mode.

```
OS10(config-if-eth1/1/1)# no switchport access vlan
```

4 Assign the interface to the controller.

```
OS10(config-if-eth1/1/1)# nve-controller
```

To view the controller information and the ports managed by the controller, use the `show nve controller` command.

```
OS10# show nve controller

Management IP            : 10.16.140.29/16
Gateway IP               : 55.55.5.5
Max Backoff              : 1000
Configured Controller    : 10.16.140.172:6640 ssl (connected)

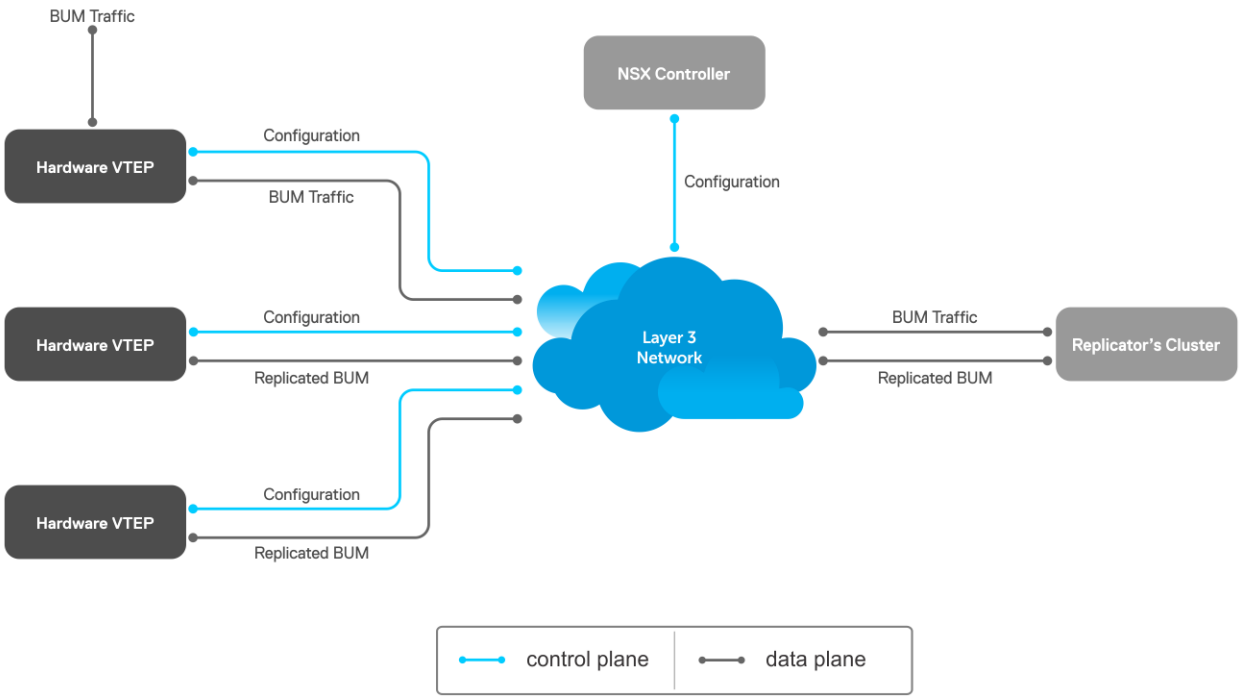
Controller Cluster
IP                Port     Protocol      Connected   State         Max-Backoff
10.16.140.173     6640     ssl           true        ACTIVE        1000
10.16.140.171     6640     ssl           false       BACKOFF       1000
10.16.140.172     6640     ssl           true        ACTIVE        1000


NVE Controller Ports
ethernet1/1/1:1
ethernet1/1/15
```

BFD in VXLAN service nodes

Service nodes (SN) handle L2 broadcast, unknown-unicast, and multicast (BUM) traffic that originates in a controller-based VXLAN environment. In SN replication, the NSX suite deploys hypervisors that replicate BUM traffic across VTEPs. These hypervisors are called *service node replicators*, which are deployed as a cluster. The NSX controller notifies the OVSDB server in the OS10 VTEP about the replication cluster through the OVSXB tables. A service node is also a VTEP, but the controller fully manages the service node. The SN is the central replication engine for flooded packets. In VXLAN, the BUM traffic replicates to all the other remote VTEPs in the same VXLAN.

The following shows BUM traffic replication in the controller environment:



A source VTEP forwards BUM traffic to the replicator, which in turn replicates the frames and transmits them to all other VTEPs through the underlay IP route.

You can use BFD to detect the health of the service nodes. BFD runs between the VTEP and the service nodes. The NSX controller provides parameters, such as the minimum TX and RX interval, and the multiplier, to initiate the BFD session. To establish a BFD session, enable the BFD on both the controller and the VTEP.

NOTE: When establishing a BFD session, the controller-provided parameters are used instead of the VTEP-configured parameters.

NOTE: If you do not enable BFD, the controller does not trigger BFD session updates and the choice of the service node does not depend on the BFD state.

If BFD is disabled in the VTEP and enabled at the NSX controller, or if BFD is disabled in the VTEP after being enabled, the service node is selected based on next-hop reachability. After BFD is enabled in the switch, a new service node with BFD enabled is selected from the list.

To view established sessions, use the `show bfd neighbors` command.

```
OS10# show bfd neighbors
* - Active session role
-------------------------------------------------------------------------------------------
  LocalAddr       RemoteAddr     Interface         State RxInt TxInt Mult VRF      Clients
-------------------------------------------------------------------------------------------
* 55.55.5.5       2.2.2.2        virtual-network0 up    1000  1000  3    default vxlan
* 55.55.5.5       2.2.2.3        virtual-network0 up    1000  1000  3    default vxlan
```

Viewing the replicators and their state

To view the state of the replicators, use the `show nve replicators` command.

- Example output without VNID

```
OS10# show nve replicators
Codes: * - Active Replicator

BFD Status:Enabled
```

```
Replicators      State
-----------------------
2.2.2.3          Up
2.2.2.2          Up
```

- Example output with VNID

```
OS10# show nve replicators vnid 10009
Codes: * - Active Replicator

BFD Status:Enabled
Replicators      State
-----------------------
2.2.2.3          Up
2.2.2.2*         Up
```

Configure and control VXLAN from VMware vCenter

You can configure and control VXLAN from the VMware vCenter Graphical User Interface (GUI). The following steps are performed in the NSX controller GUI:

1 On an OS10 switch, generate an SSL certificate in CONFIGURATION mode.

```
OS10# nve controller ssl-key-generate
```

Verify or view the certificate by using the `show nve controller ssl-certificate` command.

```
OS10# show nve controller ssl-certificate
-----BEGIN CERTIFICATE-----
MIIDgDCCAmgCAQMwDQYJKoZIhvcNAQENBQAwgYExCzAJBgNVBAYTAlVTMQswCQYD
VQQIDAJDQTEVMBMGA1UECgwMT3BlbiB2U3dpdGNoMREwDwYDVQQLDAhzd2l0Y2hj
YTE7MDkGA1UEAwwyT1ZTIHN3aXRjaGNhIENBIENlcnRpZmljYXRlICgyMDE4IFNl
cCAyMyAwMzo0NzoyMCkwHhcNMTgwOTI0MTYzMDUyWhcNMjgwOTIxMTYzMDUyWjCB
iTELMAkGA1UEBhMCVVMxCzAJBgNVBAgMAkNBMRUwEwYDVQQKDAxPcGVuIHZTd2l0
Y2gxHzAdBgNVBAsMFk9wZW4gdlN3aXRjaCBjZXJ0aWZpZXIxNTAzBgNVBAMMLGRl
bGwgaWQ6MGVlZmUwYWMtNGJjOC00MmVmLTkzOTEtN2RlMmMwY2JmMTJjMIIBIjAN
BgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAsMlD4c4fWwy+5t6VScjizlkFsNzE
BOK5PJyI3B6ReRK/J14Fdxio1YmzG0YObjxiwjpUYEsqPL3Nvh0f10KMqwqJVBdf
6sXWHUVw+9A7cIfRh0aRI+HIYyUC4YD48GlnVnaCqhxYaA0tcMzJm4r2k7AjwJUl
0pDXiqS3uJwGmfxlhvmFio8EeHM/Z79DkBRD6FUMwacAnb3yCIKZH50AWq7qRmmG
NZOgYUT+8oaj5tO/hEQfDYuv32E5z4d3FhiBJMFT86T4YvpJYyJkiKmaQWInkthL
V3VxEMXI5vJQclMhwYbKfPB4hh3+qdS5o+uVco76CVrcWi7rO3XmsBkbnQIDAQAB
MA0GCSqGSIb3DQEBDQUAA4IBAQATuFVD20GcHD8zdpYf0YaP4b6TuonUzF0jwoV+
Qr9b4kOjEBGuoPdevX3AeV/dvAa2Q6o1iOBM5z74NgHizhr067pFP841Nv7DAVb7
cPHHSSTTSeeJjIVMh0kv0KkVefsYuI4r1jqJxu0GZgBinqehXxVKlceouLvwbhb1
MFYXN3lcE2AXR746q1VIc6stNkxf3nrlOpSDz3P4VOnbAnIrY+SvUVmAT0tdrowH
99y2AzoAxUHOdWsH8EjCFch7VilmCVVhyghXdfyl6lv/F6vMRwjc343BpBW3QsGj
68ROX0ILrtOz/2q5oUb/rpJd15KFFN3itT/xYBfZ1ZdLYd5F
-----END CERTIFICATE-----
```

2 Create a VXLAN gateway.

Open a browser window, enter the vCenter IP address, and log in to VMware vCenter. To create a service node, provide the IP address and SSL certificate. Select **Home** > **Networking and Security** > **Service Definition** > **Hardware Devices**. Under the **Hardware Devices** tab, click the **Add (+)** button. The **Add Hardware Device** dialog window opens. Enter a name and copy the certificate generated in the OS10 switch and paste it in the **Certificate** box and click **OK**. This step configures the L2 gateway switch for the VXLAN virtual network.

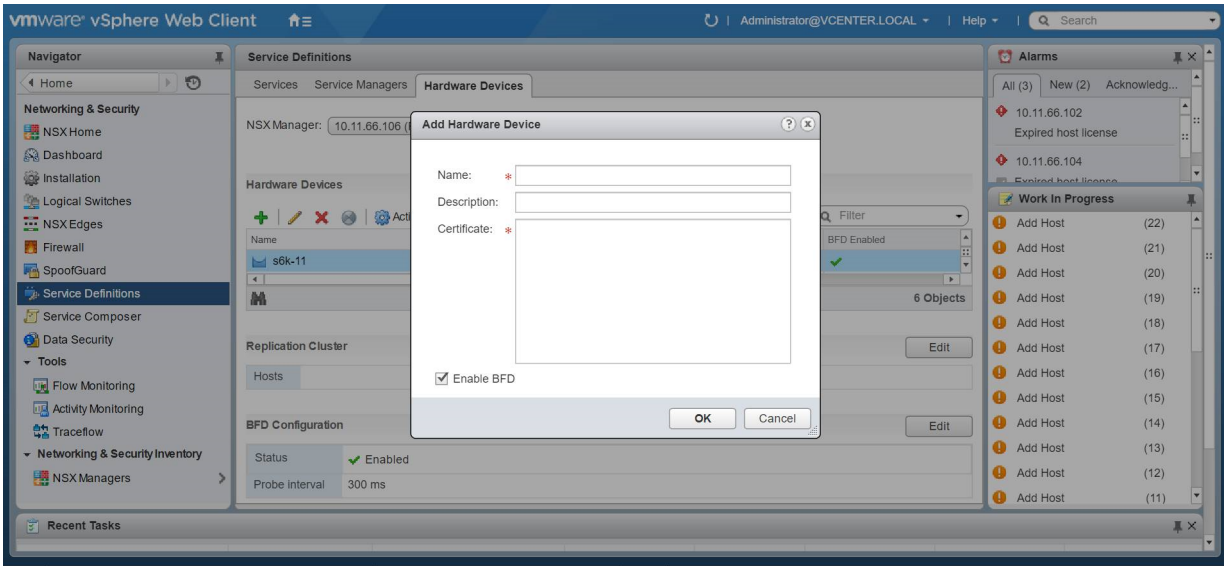
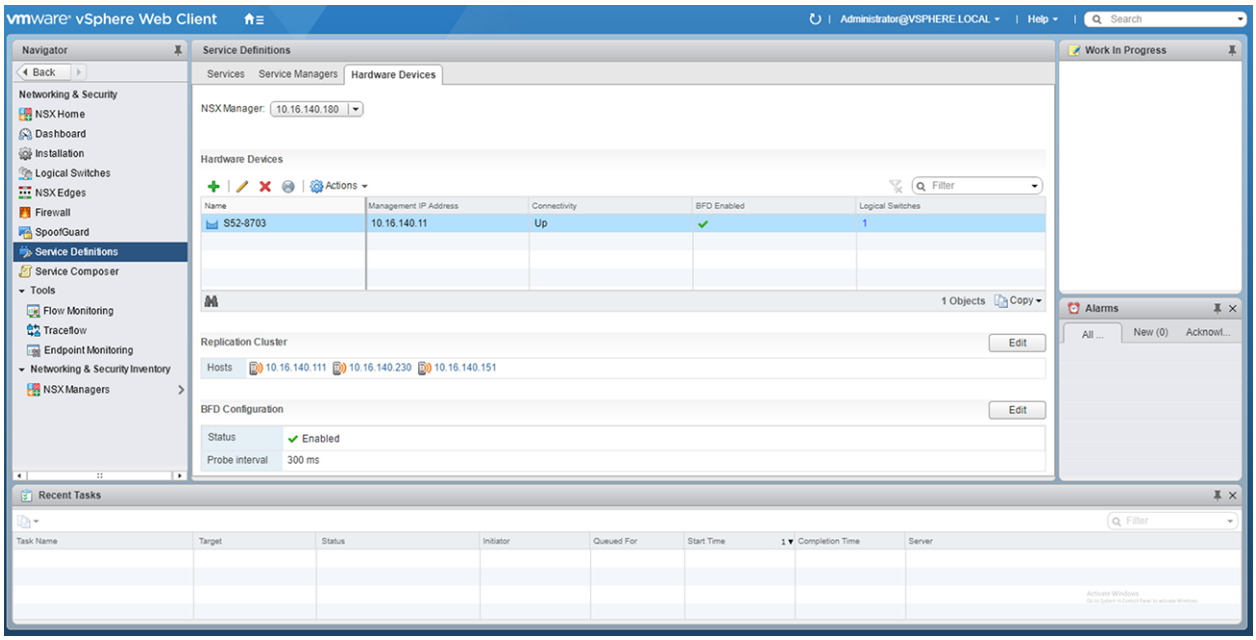


Figure 10. Add hardware device

To create a VXLAN L2 Gateway, you need the IP address of the Gateway. After connectivity establishes between the VTEP and the NSX controller, the management IP address and the connectivity status populate.



3 Create a logical switch.

You can create a logical network that acts as the forwarding domain for virtualized and nonvirtualized server workloads on the physical and virtual infrastructure.

Select **Home** > **Logical Switches** > **New Logical Switch** and click **Add**. A New Logical Switch dialog window opens. Enter a name and select **Unicast** as the replicate mode and click **OK**. This step configures the logical switch for NSX controller management.

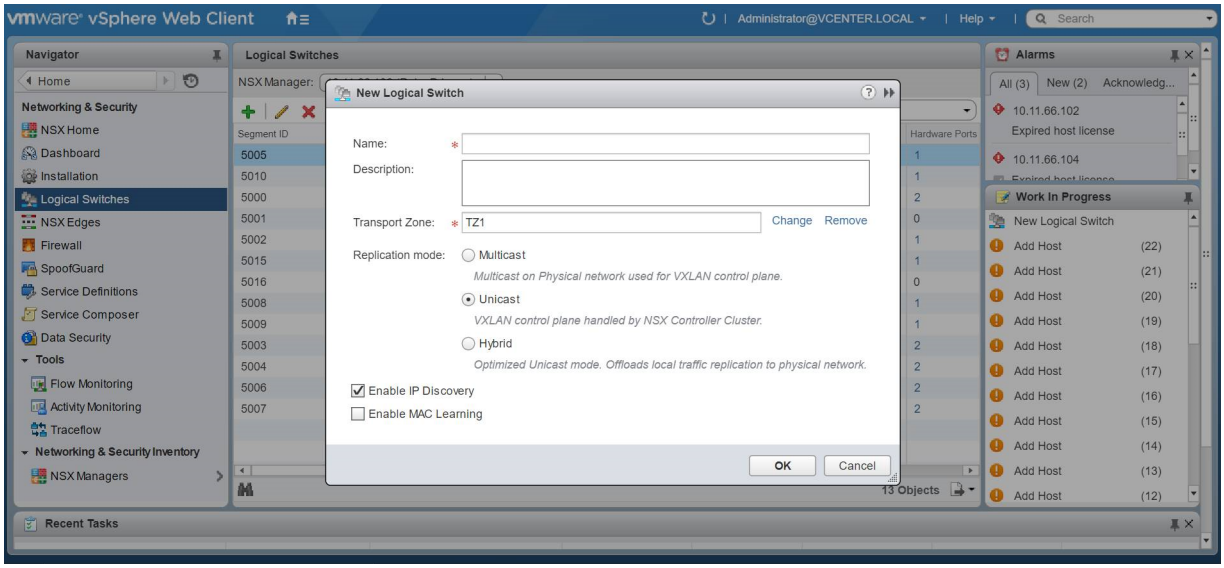
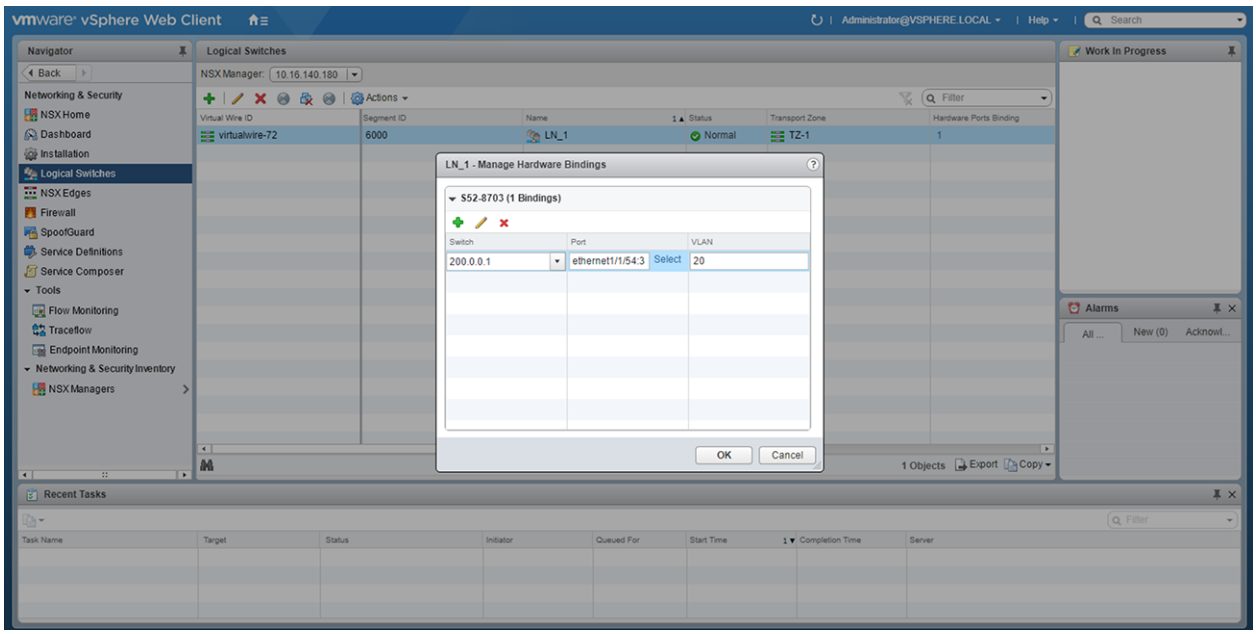


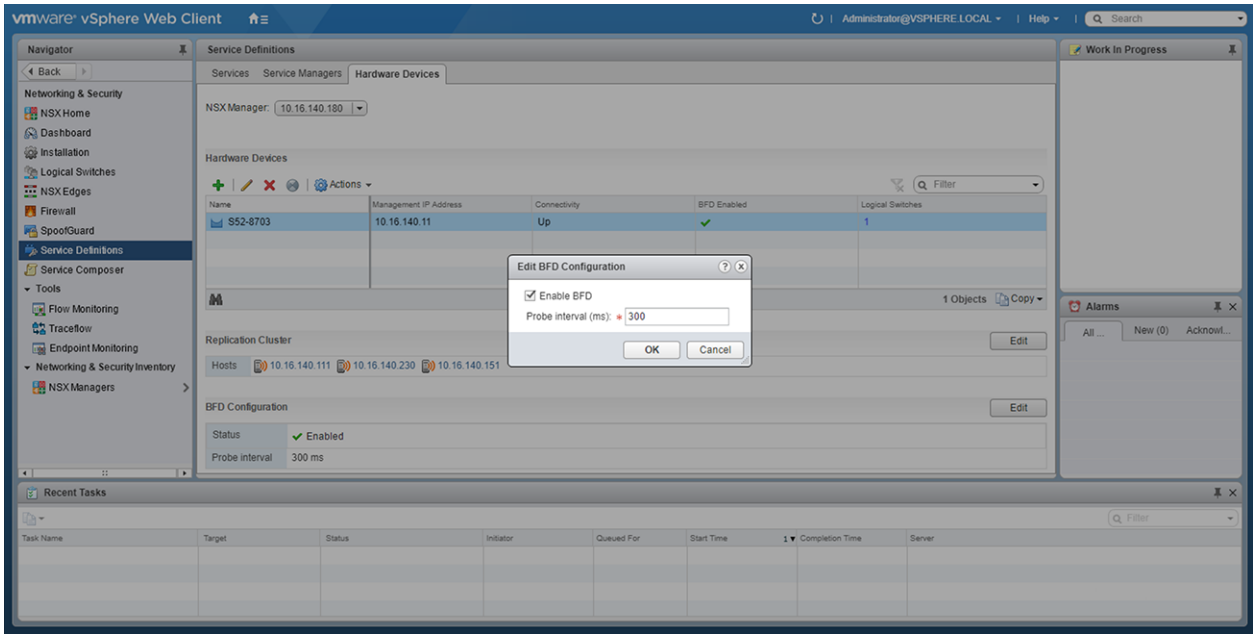
Figure 11. New logical switch

4 Create a logical switch port that provides a logical connection point for a VM interface (VIF) and a L2 gateway connection to an external network. The logical switch port binds the virtual access ports in the gateway to the logical VXLAN network and server VLAN. Select the **Logical Switches** and from the Manage Hardware Bindings dialog window, expand a VTEP and click **Add**. Select the switch from the Switch column. Enter the hardware port under the port field, and enter the VLAN ID under the VLAN column and click **OK**.



5 (Optional) Enable or disable BFD globally.

Go to **Hardware Devices** > **BFD Configuration**, and click the **Edit** button. The Edit BFD Configuration windows opens. Check or clear the **Enable BFD** check box. You can also change the probe interval if required.

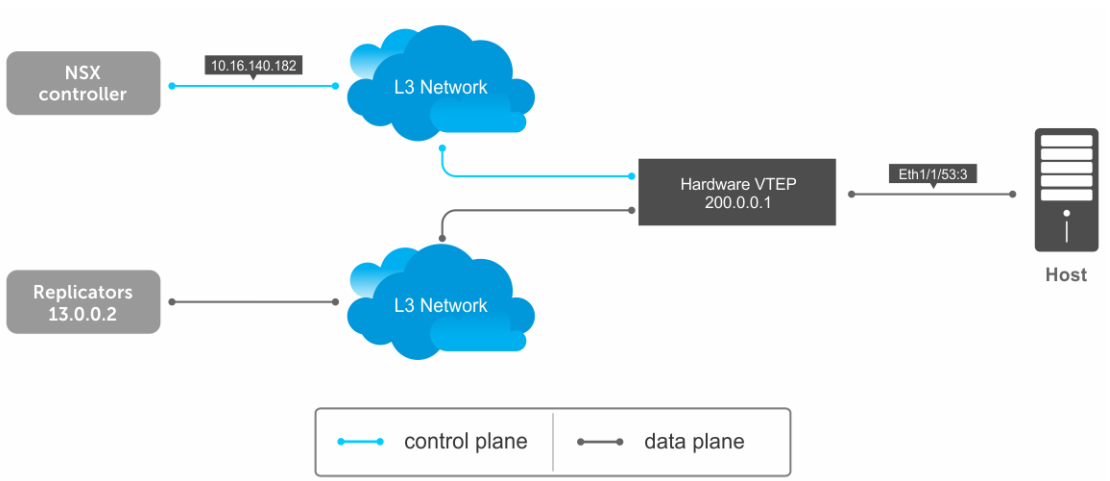


After you configure a VMware NSX controller on a server VM, connect to the controller from the VXLAN gateway switch.

For more information about the NSX controller configuration in the VTEP, see Configure a connection to an OVSDB controller. For more information about NSX controller configuration, see the *NSX User Guide* from VMware.

Example: VXLAN with a controller configuration

This example shows a simple NSX controller and an OS10 VTEP deployed in VXLAN environment.



To configure an NSX controller-based VXLAN:

- Configure the controller and an interface the controller manages, in the OS10 VTEP
- Configure the NSX controller in VMware vCenter. For more information about configuring the NSX controller using the GUI, see the Configure and control VXLAN from the VMware vCenter.

Configure the controller and interface

You must configure an OS10 VTEP with the controller configuration so that the VTEP can communicate with the NSX controller. The NSX controller handles configurations and control plane operations in the VXLAN environment.

1 Configure the underlay OSPF protocol.

```
OS10# configure terminal
OS10(config)# router ospf 1
OS10(config)# exit
OS10(config)# interface ethernet 1/1/55:1
OS10(config-if-eth1/1/55:1)# no switchport
OS10(config-if-eth1/1/55:1)# ip ospf 1 area 0.0.0.0
OS10(config-if-eth1/1/55:1)# exit
```

2 Configure a Loopback interface.

```
OS10(config)# interface loopback 1
OS10(config-if-lo-1)# no shutdown
OS10(config-if-lo-1)# ip address 200.0.0.1/32
OS10(config-if-lo-1)# exit
```

3 Create an NVE instance and configure a Loopback interface as the VXLAN source tunnel interface.

```
OS10(config)# nve
OS10(config-nve)# source-interface loopback 1
```

4 Configure an OVSDB-based controller (NSX).

```
OS10(config-nve)# controller ovsdb
OS10(config-nve-ovsdb)# ip 10.16.140.182 port 6640 ssl
OS10(config-nve-ovsdb)# max-backoff 10000
OS10(config-nve-ovsdb)# exit
```

5 Assign an interface that the controller manages.

```
OS10(config)# interface ethernet 1/1/54:3
OS10(config-if-eth1/1/54:3)# switchport mode trunk
OS10(config-if-eth1/1/54:3)# no switchport access vlan
OS10(config-if-eth1/1/54:3)# nve-controller
```

Verify the controller configuration

To view controller-based information on the VTEP, use the `show nve controller` command.

```
OS10# show nve controller

Management IP            : 10.16.140.11/16
Gateway IP               : 200.0.0.1
Max Backoff              : 10000
Configured Controller    : 10.16.140.181:6640 ssl (connected)

Controller Cluster
IP                Port     Protocol      Connected   State         Max-Backoff
10.16.140.182     6640     ssl           true        ACTIVE        10000
10.16.140.183     6640     ssl           true        ACTIVE        10000
10.16.140.181     6640     ssl           true        ACTIVE        10000

NVE Controller Ports
ethernet1/1/54:3
```

To display the VNID, port members, source interface, and remote VTEPs of the VXLAN, use the `show virtual-network` command.

```
OS10# show virtual-network
Codes: DP - MAC-learn Dataplane, CP - MAC-learn Controlplane, UUD - Unknown-Unicast-Drop
Virtual Network: 0
   Members:

Virtual Network: 6000
   Members:
      VLAN 20: ethernet1/1/54:3
   VxLAN Virtual Network Identifier: 6000
      Source Interface: loopback1(200.0.0.1)
      Remote-VTEPs (flood-list): 13.0.0.5(CP)
```

To view all the replicators and their status in the VXLAN, use the `show nve replicators` command.

```
OS10# show nve replicators
Codes: * - Active Replicator

BFD Status:Disabled
Replicators      State
-----------------------
13.0.0.5         Up
13.0.0.3         Up
13.0.0.2         Up
```

To view the remote VTEP status, use the `show nve remote-vtep` command.

```
OS10# show nve remote-vtep
IP Address: 13.0.0.2,  State: up, Encap: VxLAN
   VNI list: ,6000
IP Address: 13.0.0.3,  State: up, Encap: VxLAN
   VNI list: ,6000
IP Address: 13.0.0.5,  State: up, Encap: VxLAN
   VNI list: ,6000
```

VXLAN Controller commands

controller ovsdb

Enables Controller mode to configure the controller parameters.

Syntax	`controller ovsdb`
Parameters	None
Default	None
Command mode	CONFIGURATION-NVE
Usage information	The controller configuration initiates the OVSDB service on the OS10 switch. This populates the controller-managed ports in the OVSDB server. The `no` version of this command stops the OVSDB service. The `no` version command fails if any ports are configured as controller-managed ports or IP address configuration. **NOTE: Before removing the controller configuration from the device, you must delete all controller-managed ports and IP address configuration.**

Example

```
OS10(config)# nve
OS10(config-nve)# controller ovsdb
```

Supported releases 10.4.3.0 or later

ip port ssl

Add an OVSDB controller connection, such as IP address, port number and the connection type of session, in the switch. Currently, the only OVSDB controller supported is the NSX controller.

Syntax `ip ip-address port port-number ssl`

Parameters

- *ip-address* — Specify the IP address of the NSX controller to connect with.
- *port-number* — Specify the port number through which the connection to the NSX controller is made.

Default For an OVSDB-based controller, configure the following:

- Port number as `6640`
- Connection type as `SSL`

Command mode CONFIGURATION-NVE-OVSDB

Usage information OS10 supports only one OVSDB controller connection at a time. You can configure only one set of connection parameters for a controller. This command adds the controller configuration to OVSDB. The local OVSDB server initiates the connection to the NSX controller. The `no` version of this command removes the controller configuration from the local OVSDB server, which in turn removes the connection to the NSX controller.

Example

```
OS10(config)# nve
OS10(config-nve)# controller ovsdb
OS10(config-nve-ovsdb)# ip 10.11.66.110 port 6640 ssl
```

Supported releases 10.4.3.0 or later

max-backoff

Configures the time interval, in milliseconds, that the switch waits between the connection attempts to the controller.

Syntax `max-backoff interval`

Parameters *interval*—Enter the amount of time, in milliseconds, that the switch waits between the connection attempts to the controller, from 1000 to 180000 milliseconds.

Default 8000 milliseconds

Command Mode CONFIGURATION-NVE-OVSDB

Usage Information The `no` version of this command replaces the default maximum wait time configuration in the switch.

Example

```
OS10(config)# nve
OS10(config-nve)# controller ovsdb
OS10(config-nve-ovsdb)# max-backoff 40000
```

Supported Releases 10.4.3.0 or later

nve-controller

Assigns an controller-managed interface for VXLAN configurations.

Syntax `nve-controller`

Parameters None

Default None

Command mode INTERFACE

Usage information The interface must be in Switchport Trunk mode when adding the interface to the controller. If the interface is not in the Switchport Trunk mode, the following error message displays:

```
% Error: Interface ethernet1/1/1, must be in switchport trunk for controller
mode.
```

NOTE: If the interface has active port-scoped VLAN (Port,VLAN) pairs configured by the controller, you cannot remove an interface from the controller.

The `no` version of this command removes the interface from the controller and removes any VXLAN binding associated with the interface.

Example

```
OS10(config)# interface ethernet 1/1/1
OS10(config-if-eth1/1/1)# nve-controller
```

Supported releases 10.4.3.0 or later

nve controller ssl-key-generate

Generates the SSL certificate for the OVSDB server to setup the SSL connection with the controller.

Syntax `nve controller ssl-key-generate`

Parameters None

Default None

Command mode EXEC

Usage information This command is available only for the `sysadmin` and `secadmin` roles.

This command generates the SSL certificate and restarts the OVSDB server to pick up the newly generated certificate.

NOTE: During an OVSDB server restart, any updates to the server are missed.

Example

```
OS10# nve controller ssl-key-generate
```

Supported releases 10.4.3.0 or later

show nve controller

Displays information about the controller and the controller-managed interfaces.

Syntax `show nve controller`

Parameters None

Default None

Command mode EXEC

Example

```
OS10# show nve controller

Management IP            : 10.16.140.29/16
Gateway IP               : 55.55.5.5
Max Backoff              : 1000
Configured Controller    : 10.16.140.172:6640 ssl (connected)

Controller Cluster
IP                 Port     Protocol      Connected   State         Max-Backoff
10.16.140.173      6640     ssl           true        ACTIVE        1000
10.16.140.171      6640     ssl           false       BACKOFF       1000
10.16.140.172      6640     ssl           true        ACTIVE        1000


NVE Controller Ports
ethernet1/1/1:1
ethernet1/1/15
```

Supported releases 10.4.3.0 or later

show nve controller ssl-certificate

Displays the SSL certificate generated in the system.

Syntax `show nve controller ssl-certificate`

Parameters None

Default None

Command mode EXEC

Usage information This command is available only for `sysadmin` and `secadmin` roles.

Example

```
OS10# show nve controller
-----BEGIN CERTIFICATE-----
MIIDgDCCAmgCAQMwDQYJKoZIhvcNAQENBQAwgYExCzAJBgNVBAYTAlVTMQswCQYD
VQQIDAJDQTEVMBMGA1UECgwMT3BlbiB2U3dpdGNoMREwDwYDVQQLDAhzd2l0Y2hj
YTE7MDkGA1UEAwwyT1ZTIHN3aXRjaGNhIENBIENlcnRpZmljYXRlICgyMDE4IFNl
cCAyMyAwMzo0NzoyMCkwHhcNMTgwOTI0MTYzMDUyWhcNMjgwOTIxMTYzMDUyWjCB
iTELMAkGA1UEBhMCVVMxCzAJBgNVBAgMAkNBMRUwEwYDVQQKDAxPcGVuIHZTd2l0
Y2gxHzAdBgNVBAsMFk9wZW4gdlN3aXRjaCBjZXJ0aWZpZXIxNTAzBgNVBAMMLGRl
bGwgaWQ6MGVlZmUwYWMtNGJjOC00MmVmLTkzOTEtN2RlMmMwY2JmMTJjMIIBIjAN
BgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAsMlD4c4fWwy+5t6VScjizlkFsNzE
BOK5PJyI3B6ReRK/J14Fdxio1YmzG0YObjxiwjpUYEsqPL3Nvh0f10KMqwqJVBdf
6sXWHUVw+9A7cIfRh0aRI+HIYyUC4YD48GlnVnaCqhxYaA0tcMzJm4r2k7AjwJUl
0pDXiqS3uJwGmfxlhvmFio8EeHM/Z79DkBRD6FUMwacAnb3yCIKZH50AWq7qRmmG
NZOgYUT+8oaj5tO/hEQfDYuv32E5z4d3FhiBJMFT86T4YvpJYyJkiKmaQWInkthL
V3VxEMXI5vJQclMhwYbKfPB4hh3+qdS5o+uVco76CVrcWi7rO3XmsBkbnQIDAQAB
MA0GCSqGSIb3DQEBDQUAA4IBAQATuFVD20GcHD8zdpYf0YaP4b6TuonUzF0jwoV+
Qr9b4kOjEBGuoPdevX3AeV/dvAa2Q6o1iOBM5z74NgHizhr067pFP841Nv7DAVb7
```

```
cPHHSSTTSeeJjIVMh0kv0KkVefsYuI4r1jqJxu0GZgBinqehXxVKlceouLvwbhb1
MFYXN3lcE2AXR746q1VIc6stNkxf3nrlOpSDz3P4VOnbAnIrY+SvUVmAT0tdrowH
99y2AzoAxUHOdWsH8EjCFch7VilmCVVhyghXdfyl6lv/F6vMRwjc343BpBW3QsGj
68ROX0ILrtOz/2q5oUb/rpJd15KFFN3itT/xYBfZ1ZdLYd5F
-----END CERTIFICATE-----
```

Supported releases 10.4.3.0 or later

show nve replicators

Displays all the replicators and their states.

Syntax `show nve replicators [vnid vnid]`

Parameters None

Default None

Command mode EXEC

Usage information When you specify the VNID, the output displays the active replicator selected for that VNID based on the load-balancing logic.

Example (without VNID)

```
OS10# show nve replicators
Codes: * - Active Replicator

BFD Status:Enabled
Replicators      State
-----------------------
2.2.2.3          Up
2.2.2.2          Up

OS10# show nve replicators
```

Example (with VNID)

```
OS10# show nve replicators vnid 10009
Codes: * - Active Replicator

BFD Status:Enabled
Replicators      State
-----------------------
2.2.2.3          Up
2.2.2.2*         Up
```

Supported releases 10.4.3.0 or later

show ovsdb-tables mac-local-ucast

Displays information about local MAC address entries including each MAC address, IP address, local switch name, and VNID.

Syntax `show ovsdb-tables mac-local-ucast`

Parameters None

Default None

Command mode EXEC

Usage information This command is available only for `netadmin`, `sysadmin`, and `secadmin` roles.

Example

```
OS10# show ovsdb-tables mac-local-ucast
Count : 1356
Ucast_Macs_Local table
```

```
MAC                _uuid                                ipaddr locator
logical_switch
------------------ ------------------------------------ ------ ------------------------------------
---------------------
"00:00:09:00:00:00" 948d2357-9a68-49b2-b5b2-a6a9beaec17a ""     bb43d2ec-1e60-4367-9840-648a8cc8acff f8994210-
e29d-4ad4-90fb-557c30f83769
"00:00:09:00:00:01" 4e620093-311a-420e-957f-fbd2bb63f20a ""     bb43d2ec-1e60-4367-9840-648a8cc8acff f8994210-
e29d-4ad4-90fb-557c30f83769
"00:00:09:00:00:02" 3846973c-2b29-4c84-af39-dfe7513cdb3d ""     bb43d2ec-1e60-4367-9840-648a8cc8acff f8994210-
e29d-4ad4-90fb-557c30f83769
```

Supported releases 10.4.3.0 or later

show ovsdb-tables mac-remote-ucast

Displays information about remote MAC address entries including each MAC address, IP address, local switch name, and VNID.

Syntax `show ovsdb-tables mac-remote-ucast`

Parameters None

Default None

Command mode EXEC

Usage information This command is available only for `netadmin`, `sysadmin`, and `secadmin` roles.

Example

```
OS10# show ovsdb-tables mac-remote-ucast
Count : 1
Ucast_Macs_Remote table
MAC                _uuid                                ipaddr locator
logical_switch
------------------ ------------------------------------ ------ ------------------------------------
----------------------"00:50:56:8a:b4:c8" 61fa240b-e6a3-4d8e-a693-dd2468e6f308 ""     3105e34b-a273-4193-
a60f-51d9cee91403 6932fc02-fb12-4a22-9ec2-f0e2b20df476
```

Supported releases 10.4.3.0 or later

show ovsdb-tables manager

Displays information about each OVSDB manager and parameters including connection status and a list of physical switch names.

Syntax `show ovsdb-tables manager`

Parameters None

Default None

Command mode EXEC

Usage information This command is available only for `netadmin`, `sysadmin`, and `secadmin` roles.

Example

```
OS10# show ovsdb-tables manager
Count : 3
Manager table
_uuid                                inactivity_probe is_connected max_backoff
other_config
status                                                                  target
------------------------------------ ---------------- ------------ -----------
------------ ------------------------------
478ec8ca-9c5a-4d29-9069-633af6c48002 []               false        1000
{}
{state=BACKOFF}                                                         "ssl:
10.16.140.171:6640"
52f2b491-6372-43e0-98ed-5c4ab0ca8542 []               true         1000
{}           {sec_since_connect="37831", sec_since_disconnect="37832",
state=ACTIVE} "ssl:10.16.140.173:6640"
```

```
7b8a7e36-6221-4297-b85e-51f910abcb5c []                 true          1000
{}           {sec_since_connect="87", sec_since_disconnect="99",
state=ACTIVE}       "ssl:10.16.140.172:6640"
OS10#
```

Supported releases 10.4.3.0 or later

show ovsdb-tables tunnel

Displays information about the tunnels the physical switch creates.

Syntax `show ovsdb-tables tunnel`

Parameters None

Default None

Command mode EXEC

Usage information This command is available only for `netadmin`, `sysadmin`, and `secadmin` roles.

Example

```
OS10# show ovsdb-tables tunnel
Count : 2
Tunnel table
_uuid                                bfd_config_local
bfd_config_remote                                    bfd_params
bfd_status                                                                    local
remote
------------------------------------ ------------------------------------------------------------ ------------------------------8025d953-
acf5-4091-9fa2-75d41953b397 {bfd_dst_ip="55.55.5.5", bfd_dst_mac="00:23:20:00:00:01"} {bfd_dst_ip="2.2.2.2",
bfd_dst_mac="00:50:56:65:b2:3c"} {enable="true", forwarding_if_rx="true", min_rx="1000"} {diagnostic="No Diagnostic",
enabled="true", forwarding="true", remote_state=up, state=up} bb43d2ec-1e60-4367-9840-648a8cc8acff 2d8963da-24d0-4fbd-81e2-
fb1a7bba88fd
9853f77a-9db7-47f5-8203-b5b8895d15bd {bfd_dst_ip="55.55.5.5", bfd_dst_mac="00:23:20:00:00:01"} {bfd_dst_ip="2.2.2.3",
bfd_dst_mac="00:50:56:6e:56:9b"} {enable="true", forwarding_if_rx="true", min_rx="1000"} {diagnostic="No Diagnostic",
enabled="true", forwarding="true", remote_state=up, state=up} bb43d2ec-1e60-4367-9840-648a8cc8acff 5eee586b-
e0aa-442b-83ea-16633ec41230
```

Supported releases 10.4.3.0 or later

8

UFT modes

A switch in a Layer 2 (L2) network may require a larger MAC address table size, while a switch in a Layer 3 (L3) network may require a larger routing table size. Unified forwarding table (UFT) offers the flexibility to configure internal L2/L3 forwarding table sizes.

OS10 supports several UFT modes for the forwarding tables. By default, OS10 selects a UFT mode that provides a reasonable size for all tables. The supported UFT modes are: default, scaled-l2–switch, scaled-l3–hosts, and scaled-l3–routes.

NOTE: S5148–ON does not support UFT modes.

Table 24. UFT Modes — Table Size for S4048–ON, S4048T-ON, S6010–ON

UFT Mode	L2 MAC Table Size	L3 Host Table Size	L3 Routes Table Size
Scaled-l2–switch	294912	16384	16384
Scaled-l3–hosts	98304	212992	98304
Scaled-l3–routes	32768	16384	131072
Default	163840	147456	16384

Table 25. UFT Modes — Table Size for S3048–ON

UFT Mode	L2 MAC Table Size	L3 Host Table Size	L3 Routes Table Size
Scaled-l2–switch	40960	2048	8192
Scaled-l3–hosts	8192	18432	8192
Default	28672	8192	8192

Table 26. UFT Modes — Table Size for S41XX-ON series

UFT Mode	L2 MAC Table Size	L3 Host Table Size	L3 Routes Table Size
Scaled-l2–switch	278528	4096	16384
Scaled-l3–hosts	16384	266240	16384
Scaled-l3–routes	16384	4096	262144
Default	81920	69632	131072

Table 27. UFT Modes — Table Size for Z9100-ON

UFT Mode	L2 MAC Table Size	L3 Host Table Size	L3 Routes Table Size
Scaled-l2–switch	139264	8192	16384
Scaled-l3–hosts	8192	139264	16384
Scaled-l3–routes	8192	8192	131072
Default	73728	73728	16384

Table 28. UFT Modes — Table Size for Z9264F-ON

UFT Mode	L2 MAC Table Size	L3 Host Table Size	L3 Routes Table Size
Scaled-l2–switch	270336	8192	32768
Scaled-l3–hosts	8192	270336	32768
Scaled-l3–routes	8192	8192	262144
Default	139264	139264	32768

Table 29. UFT Modes — Table Size for S52XX-ON series

UFT Mode	L2 MAC Table Size	L3 Host Table Size	L3 Routes Table Size
Scaled-l2–switch	294912	16384	16384
Scaled-l3–hosts	32768	278528	16384
Scaled-l3–routes	32768	16384	389120
Default	163840	147456	16384

Table 30. UFT Modes — Table Size for S42xxFB-ON

UFT Mode	L2 MAC Table Size	L3 Host Table Size	L3 Routes Table Size
Default	250 K	48 K	130 K

Table 31. UFT Modes — Table Size for S42xxFBL-ON

UFT Mode	L2 MAC Table Size	L3 Host Table Size	L3 Routes Table Size
Default	250 K	48 K	2 Million

NOTE: The L3 routes table size for Scaled-l3-routes mode might vary depending on the routes that are being installed.

Configure UFT modes

Available UFT modes include L2 MAC table, L3 host table, or L3 route table sizes.

- Select a mode to initialize the maximum table size in CONFIGURATION mode.
  ```
  hardware forwarding-table mode [scaled-l2 | scaled-l3-routes | scaled-l3-hosts]
  ```
- Disable UFT mode in CONFIGURATION mode.
  ```
  no hardware forwarding-table
  ```

Configure UFT mode

```
OS10(config)# hardware forwarding-table mode scaled-l3-hosts
```

View UFT mode information

```
OS10# show hardware forwarding-table mode
                       Current Settings           Next-boot Settings
 Mode                  default-mode               scaled-l3-hosts
 L2 MAC Entries    :   163840                     98304
 L3 Host Entries   :   147456                     212992
 L3 Route Entries  :   16384                      98304
```

View UFT information for all modes

```
OS10# show hardware forwarding-table mode all
 Mode                default     scaled-l2   scaled-l3-routes   scaled-l3-hosts
 L2 MAC Entries      163840      294912      32768              98304
 L3 Host Entries     147456      16384       16384              212992
 L3 Route Entries    16384       16384       131072             98304
```

IPv6 extended prefix routes

IPv6 addresses that contain prefix routes with mask between /64 to /128 are called as IPv6 extended prefix routes. These routes require double the key size in the Longest prefix match (LPM) table.

You can configure the number of route entries for extended prefix using the `hardware l3 ipv6-extended-prefix` *`prefix-number`* command.

Save and Reload the switch for the settings to become effective.

Configure IPv6 extended prefix route

```
OS10# configure terminal
OS10(config)# hardware l3 ipv6-extended-prefix 2048
 % Warning: IPv6 Extended Prefix Installation will be applied only after a save and reload.
OS10(config)# do write memory
OS10(config)# reload
```

View IPv6 extended prefix route configuration

```
OS10# show running-configuration | grep hardware
 hardware l3 ipv6-extended-prefix 2048
```

Configuration before reload:

```
OS10# show hardware l3
                              Current Settings       Next-boot Settings
 IPv6 Extended Prefix Entries: 0                      2048
```

Configuration after reload:

```
OS10# show hardware l3
                              Current Settings       Next-boot Settings
 IPv6 Extended Prefix Entries: 2048                   2048
```

The `no` version of the command removes the IPv6 extended prefix route configuration. Save and Reload the switch to remove the configuration.

```
OS10(config)# no hardware l3 ipv6-extended-prefix
 % Warning: Un-configuring IPv6 Extended Prefix will be applied only after a save and reload.
```

UFT commands

hardware forwarding-table mode

Selects a mode to initialize the maximum scalability size. The available options are: scaled L2 MAC address table, scaled L3 routes table, or scaled L3 hosts table.

Syntax `hardware forwarding-table mode {scaled-l2 | scaled-l3-routes | scaled-l3-hosts}`

Parameters

- `scaled-l2` —Enter the L2 MAC address table size.
- `scaled-l3-routes` — Enter the L3 routes table size.
- `scaled-l3-hosts` — Enter the L3 hosts table size.

Defaults The default parameters vary according to the platform. See UFT modes.

Command Mode CONFIGURATION

Usage Information Configure the sizes of internal L2 and L3 forwarding tables for your requirements of the network environment. To apply the changes, reload the switch.

The `no` version of this command resets the UFT mode to default.

Example

```
OS10(config)# hardware forwarding-table mode scaled-l3-hosts
```

Supported Releases 10.3.0E or later

hardware l3 ipv6-extended-prefix

Configures the maximum number of route entries for IPv6 extended prefix route.

Syntax `hardware l3 ipv6-extended-prefix` *`prefix-number`*

Parameters *`prefix-number`* —Enter the maximum number of route entries for IPv6 extended prefix route. The options available are: 1024, 2048, or 3072.

Defaults None

Command Mode CONFIGURATION

Usage Information Save and Reload the switch for the settings to become effective. The `no` version of the command removes the IPv6 extended prefix route configuration.

Example

```
OS10# configure terminal
OS10(config)# hardware l3 ipv6-extended-prefix 2048
```

```
 % Warning: IPv6 Extended Prefix Installation will be applied only after a
save and reload.
OS10(config)# do write memory
OS10(config)# reload
```

Supported Releases 10.4.1.0 or later

show hardware forwarding-table mode

Displays the current hardware forwarding table mode, and the mode after the next boot.

Syntax
```
show hardware forwarding-table mode
```

Parameters None

Defaults None

Command Mode EXEC

Usage Information None

Example
```
OS10# show hardware forwarding-table mode
                        Current Settings          Next-boot Settings
 Mode                   default-mode              scaled-l3-hosts
 L2 MAC Entries    :    163840                    98304
 L3 Host Entries   :    147456                    212992
 L3 Route Entries  :    16384                     98304
```

Supported Releases 10.3.0E or later

show hardware forwarding-table mode all

Displays table sizes for the hardware forwarding table modes.

Syntax
```
show hardware forwarding-table mode all
```

Parameters None

Defaults None

Command Mode EXEC

Usage Information None

Example
```
OS10# show hardware forwarding-table mode all
 Mode               default     scaled-l2   scaled-l3-routes   scaled-l3-hosts
 L2 MAC Entries     163840      294912      32768              98304
 L3 Host Entries    147456      16384       16384              212992
 L3 Route Entries   16384       16384       131072             98304
```

Supported Releases 10.3.0E or later

show hardware l3

Displays the IPv6 extended prefix route configuration.

Syntax
```
show hardware l3
```

Parameters None

Defaults None

Command Mode EXEC

Usage Information None

Example

```
OS10# show hardware l3
                                Current Settings       Next-boot Settings
 IPv6 Extended Prefix Entries: 2048                    2048
```

Supported Releases 10.4.1.0 or later

9

System management

Dynamic Host Configuration Protocol	Provides information to dynamically assign IP addresses and other configuration parameters to network hosts based on policies, see DHCP commands.
Network Time Protocol	Provides information to synchronize timekeeping between time servers and clients, see NTP commands.
RESTCONF API	Enables you to configure and monitor an OS10 switch using JSON-structured messages to access YANG-modeled data in configuration data stores, see CLI commands for RESTCONF API.
Security	Provides information about role-based access control, RADIUS server, user roles, and user names, see Security commands.
Simple Network Management Protocol	Provides a message format for communication between Simple Network Management Protocol (SNMP) managers and agents. SNMP provides a standardized framework and common language for network monitoring and device management, see SNMP commands.
OS10 image upgrade	Provides information to upgrade the OS10 software image, see Upgrade commands.

Dynamic Host Configuration Protocol

Dynamic Host Configuration Protocol (DHCP) is an application layer protocol that dynamically assigns IP addresses and other configuration parameters to network end-stations, also known as hosts, based on configuration policies network administrators determine.

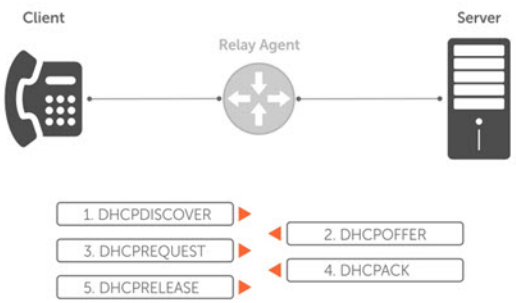


Figure 12. Client and Server Messaging

DHCP server	Network device offering configuration parameters to the client.
DHCP client	Network device requesting configuration parameters from the server.
Relay agent	Intermediary network device that passes DHCP messages between the client and the server when the server is not on the same subnet as the host.

Packet format and options

The DHCP server listens on port 67 and transmits to port 68. The DHCP client listens on port 68 and transmits to port 67.

Configuration parameters are options in the DHCP packet in type, length, value (TLV) format. To limit the number of parameters that servers must provide, hosts enter the parameters that they require and the server sends only those parameters. DHCP uses the User Datagram Protocol (UDP) as its transport protocol.

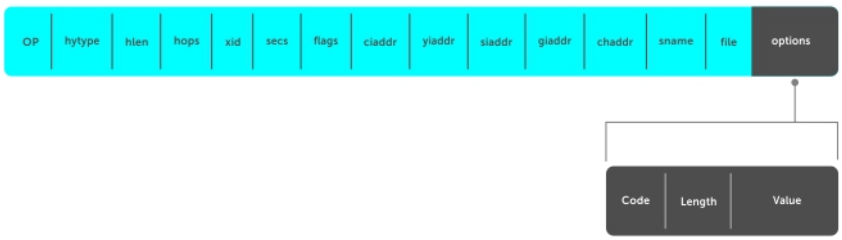


Figure 13. DHCP Packet Format

The table shows common options using DHCP packet formats.

DHCP Option	Description
Subnet mask	1 — Subnet mask of the client
Router	3 — Router IP addresses that serve as the default gateway for the client
Domain name server	6 — Domain name servers (DNS) that are available to the client
Domain name	15 — Domain name that clients use to resolve hostnames via DNS
IP address lease time	51 — Amount of time that the client uses an assigned IP address
DHCP message type	53: • 1 — DHCPDISCOVER • 2 — DHCPOFFER • 3 — DHCPREQUEST • 4 — DHCPDECLINE • 5 — DHCPACK • 6 — DHCPNACK • 7 — DHCPRELEASE • 8 — DHCPINFORM
Parameter request list	55 — Parameters the server requires for DHCP clients. This is a series of octets where each octet is a DHCP option code
Renewal time	58 — Amount of time, after the IP address is granted, that the client attempts to renew its lease with the *original* server
Rebinding time	59 — Amount of time, after the IP address is granted, that the client attempts to renew its lease with *any* server, if the original server does not respond
Vendor class identifier	60 — User-defined string the Relay Agent uses to forward DHCP client packets to a specific DHCP server
User port stacking	230 — Stacking option variable that provides the DHCP server stack-port details when the DHCP offer is set
End	255 — Signal of the last option in the DHCP packet

DHCP server

The Dynamic Host Configuration Protocol (DHCP) server provides network configuration parameters to DHCP clients on request. A DHCP server dynamically allocates four required IP parameters to each computer on the virtual local area network (VLAN) — the IP address, network mask, default gateway, and name server address. DHCP IP address allocation works on a client/server model where the server assigns the client reusable IP information from an address pool.

DHCP automates network-parameter assignment to network devices. Even in small networks, DHCP makes it easier to add new devices to the network. The DHCP access service provides a centralized, server-based setup to add clients to the network. This setup means you do not have to manually create and maintain IP address assignments for clients.

When you use DHCP to manage a pool or IP addresses among hosts, you reduce the number of IP addresses you need. DHCP manages the IP address pool by leasing an IP address to a host for a limited period, allowing the DHCP server to share a limited number of IP addresses. DHCP also provides a central database of devices that connects to the network and eliminates duplicate resource assignments.

Automatic address allocation

Automatic address allocation is an address assignment method that the DHCP server uses to lease an IP address to a client from a pool of available addresses. You cannot configure an empty DHCP pool, under a DHCP pool configuration. For a successful commit, you must have either a network statement or host/hardware-address (manual binding) configuration. An IP address pool is a range of addresses that the DHCP server assigns. The subnet number indexes the address pools.

1 Enable DHCP server-assigned dynamic addresses on an interface in DHCP <*POOL*> mode.

```
ip dhcp server
```

2 Create an IP address pool and provide a name in DHCP mode.

```
pool name
```

3 Enter the subnet from which the DHCP server may assign addresses in DHCP <*POOL*> mode. The `network` option specifies the subnet address. The `prefix-length` option specifies the number of bits used for the network portion of the address (18 to 31).

```
network network/prefix-length
```

4 Enter a range of IP addresses from the subnet specified above, which the DHCP server uses to assign addresses in DHCP <*POOL*> mode.

```
range {ip-address1 [ip-address2]}
```

NOTE: NOTE: Configure at least one interface to match one of the configured network pools. An interface matches a network pool when its IP address is included in the subnet defined for that network pool. For example, an interface with IP address 10.1.1.1/24 matches a pool configured with network 10.1.1.0/24.

DHCP server automatic address allocation

```
OS10(config)# ip dhcp server
OS10(config-dhcp)# pool Dell
OS10(config-dhcp-Dell)# default-router 20.1.1.1
OS10(config-dhcp-Dell)# network 20.1.1.0/24
OS10(config-dhcp-Dell)# range 20.1.1.2 20.1.1.8
```

Show running configuration

```
OS10(conf-dhcp-Dell)# do show running-configuration
...
!
ip dhcp server
 !
 pool Dell
  network 20.1.1.0/24
  default-router 20.1.1.1
  range 20.1.1.2 20.1.1.8
```

Address lease time

Use the `lease {days [hours] [minutes] | infinite}` command to configure an address lease time (default 24 hours).

```
OS10(config)# ip dhcp server
OS10(conf-dhcp)# pool Dell
OS10(conf-dhcp-Dell)# lease 36
```

Default gateway

Ensure the IP address of the default router is on the same subnet as the client.

1 Enable DHCP server-assigned dynamic addresses on an interface in CONFIGURATION mode.

```
ip dhcp server
```

2 Create an IP address pool and provide a name in DHCP mode.

```
pool name
```

3 Enter the default gateway(s) for the clients on the subnet in order of preference in DHCP<*POOL*> mode.

```
default-router address
```

Change default gateway name

```
OS10(config)# ip dhcp server
OS10(conf-dhcp)# pool Dell
OS10(conf-dhcp-Dell)# default-router 20.1.1.1
```

Enable DHCP server

Use the `ip dhcp server` command to enable DHCP server-assigned dynamic addresses on an interface in CONFIGURATION mode. The DHCP server is disabled by default.

```
OS10(config)# ip dhcp server
OS10(conf-dhcp)# no disable
```

Hostname resolution

You have two choices for hostname resolution — domain name server (DNS) or NetBIOS Windows internet naming service (WINS). Both DHCP and WINS clients query IP servers to compare host names to IP addresses.

1 Enable DHCP server-assigned dynamic addresses on an interface in DHCP <*POOL*> mode.

```
ip dhcp server
```

2 Create in IP address pool and enter the name in DHCP mode.

```
pool name
```

3 Create a domain and enter the domain name in DHCP <*POOL*> mode.

```
domain-name name
```

4 Enter the DNS servers in order of preference that are available to a DHCP client in DHCP <*POOL*> mode.

```
dns-server address
```

DNS address resolution

```
OS10(config)# ip dhcp server
OS10(conf-dhcp)# pool Dell
OS10(conf-dhcp-Dell)# domain-name dell.com
OS10(conf-dhcp-Dell)# dns-server 192.168.1.1
```

NetBIOS WINS address resolution

DHCP clients can be one of four types of NetBIOS nodes — broadcast, peer-to-peer, mixed, or hybrid. Dell EMC recommends using hybrid as the NetBIOS node type.

1 Enable DHCP server-assigned dynamic addresses on an interface in DHCP <*POOL*> mode.

```
ip dhcp server
```

2 Create an IP address pool and enter the pool name in DHCP mode.

```
pool name
```

3 Enter the NetBIOS WINS name servers in order of preference that are available to DHCP clients in DHCP <*POOL*> mode.

```
netbios-name-server ip-address
```

4 Enter the keyword `Hybrid` as the NetBIOS node type in DHCP <*POOL*> mode.

```
netbios-node-type type
```

Configure NetBIOS WINS address resolution

```
OS10(config)# ip dhcp server
OS10(conf-dhcp)# pool Dell
OS10(conf-dhcp-Dell)# netbios-name-server 192.168.10.5
OS10(conf-dhcp-Dell)# netbios-node-type Hybrid
```

Manual binding entries

Address binding is mapping between the IP address and the media access control (MAC) address of a client. The DHCP server assigns the client an available IP address automatically and then creates an entry in the binding table. You can also manually create an entry for a client. Manual bindings help to guarantee that a particular network device receives a particular IP address.

Consider manual bindings as single-host address pools. There is no limit to the number of manual bindings, but you can only configure one manual binding per host. Manual binding entries do not display in the `show ip dhcp binding` output.

1 Create an address pool in DHCP mode.

```
pool name
```

2 Enter the client IP address in DHCP <*POOL*> mode.

```
host address
```

3 Enter the client hardware address in DHCP <*POOL*> mode.

```
hardware-address hardware-address
```

Configure manual binding

```
OS10(config)# ip dhcp server
OS10(conf-dhcp)# pool static
OS10(conf-dhcp-static)# host 20.1.1.2
OS10(conf-dhcp-static)# hardware-address 00:01:e8:8c:4d:0a
```

View DHCP binding table

```
OS10# show ip dhcp binding
  IP Address        Hardware address     Lease expiration        Hostname
+-------------------------------------------------------------------------
11.1.1.254        00:00:12:12:12:12 Jan 27 2016 06:23:45

Total Number of Entries in the Table = 1
```

With a fixed host configuration, also known as manual binding, you must configure a network pool with a matching subnet. The static host-to-MAC address mapping pool inherits the network mask from the network pool with subnet configuration, which includes the host's address range.

Consider the following example:

```
OS10# show running-configuration interface ethernet 1/1/2
!
interface ethernet1/1/2
no shutdown
no switchport
ip address 100.1.1.1/24
flowcontrol receive off

OS10# show running-configuration ip dhcp
!
ip dhcp server
no disable
!
pool host1
host 100.1.1.34
hardware-address 00:0c:29:ee:4c:f4
!
pool hostnetwork
lease infinite
network 100.1.1.0/24
!
pool host2
host 20.1.1.34
hardware-address 00:0c:29:aa:22:f4
```

In this example, the pool `host1`, which is the fixed host mapping pool, inherits the subnet and other attributes from the pool `hostnetwork`, which is the DHCP client IP address pool. There is no matching network pool for `host2`. Therefore, the DHCP client with the MAC address `00:0c:29:aa:22:f4` does not obtain the correct parameters.

DHCP relay agent

A DHCP relay agent relays DHCP messages to and from a remote DHCP server, even if the client and server are on different IP networks. You can configure the IP address of the remote DHCP server.

You can configure a device either as a DHCP server or a DHCP relay agent — but not both.

The DHCP relay agent supports multi-virtual routing and forwarding (VRF) instances. The client-facing and server-facing interfaces must be in the same VRF.

The DHCPv6 relay agent performs the same role as that of a DHCP relay agent, but in an IPv6 network. The DHCP relay agent forwards the DHCPv4/DHCPv6 messages from the configured interface to the DHCPv6 server as a unicast message. The DHCP relay agent then forwards the server's response to the client.

When you configure DHCPv6 relay on an interface, you must:

- Configure an IPv6 address on the interface.
- Ensure that the DHCPv6 server is reachable.

Option 82 for security

DHCP, as defined by RFC 2131, provides no authentication or security mechanisms. To provide security, the DHCP relay agent supports Option-82 with Circuit ID sub option, which is the printable name of the interface where the client request was received.

This option secures all DHCP traffic that goes through a DHCP relay agent, and ensures that communication between the DHCP relay agent and the DHCP server is not compromised.

The DHCP relay agent inserts Option 82 before forwarding DHCP packets to the DHCP server. The DHCP server includes Option 82 back in its response to the relay agent. The relay agent uses this information to forward a reply out the interface on which the request was received rather than flooding it on the entire VLAN. However, the relay agent removes Option 82 from its DHCP responses before forwarding the responses to the client.

NOTE: Option 82 is supported, but not configurable.

View DHCP Information

Use the `show ip dhcp binding` command to view the DHCP binding table entries.

View DHCP Binding Table

```
OS10# show ip dhcp binding
  IP Address        Hardware address     Lease expiration        Hostname
+-------------------------------------------------------------------------
11.1.1.254        00:00:12:12:12:12 Jan 27 2016 06:23:45

Total Number of Entries in the Table = 1
```

System domain name and list

If you enter a partial domain, the system searches different domains to finish or fully qualify that partial domain. A fully qualified domain name (FQDN) is any name that terminates with a period or dot.

OS10 searches the host table first to resolve the partial domain. The host table contains both statically configured and dynamically learned host and IP addresses. If OS10 cannot resolve the domain, it tries the domain name assigned to the local system. If that does not resolve the partial domain, the system searches the list of domains configured.

You can configure the `ip domain-list` command up to five times to enter a list of possible domain names. The system searches the domain names in the order they were configured until a match is found or the list is exhausted.

1 Enter a domain name in CONFIGURATION mode (up to 64 alphanumeric characters).

```
ip domain-name name
```

2 Add names to complete unqualified host names in CONFIGURATION mode.

```
ip domain-list name
```

Configure local system domain name and list

```
OS10(config)# ip domain-name ntengg.com
OS10(config)# ip domain-list dns1
OS10(config)# ip domain-list dns2
OS10(config)# ip domain-list dns3
OS10(config)# ip domain-list dns4
OS10(config)# ip domain-list dns5
```

View local system domain name information

```
OS10# show running-configuration

! Version 10.2.9999E
! Last configuration change at Feb  20 04:50:33 2017
!
username admin password $6$q9QBeYjZ$jfxzVqGhkxX3smxJSH9DDz7/3OJc6m5wjF8nnLD7/VKx8SloIhp4NoGZs0I/
UNwh8WVuxwfd9q4pWIgNs5BKH.
```

```
aaa authentication system:local
ip domain-name dell.com
ip domain-list f10.com
ip name-server 1.1.1.1 2::2
ip host dell-f10.com 10.10.10.10
snmp-server community public read-only
snmp-server contact http://www.dell.com/support/
snmp-server location United States
debug radius false
```

DHCP commands

default-router address

Assigns a default gateway to clients based on the IP address pool.

Syntax	`default-router address [address2...address8]`
Parameters	• *address* — Enter an IPv4 or IPv6 address to use as the default gateway for clients on the subnet in A.B.C.D or A::B format. • *address2...address8* — (Optional) Enter up to eight IP addresses, in order of preference.
Default	Not configured
Command Mode	DHCP-POOL
Usage Information	Configure up to eight IP addresses, in order of preference. Use the `no` version of this command to remove the configuration.
Example	`OS10(conf-dhcp-20.1.1.1)# default-router 20.1.1.100`
Supported Releases	10.2.0E or later

disable

Disables the DHCP server.

Syntax	`disable`
Parameters	None
Default	Disabled
Command Mode	DHCP
Usage Information	The `no` version of this command enables the DHCP server.
Example	`OS10(conf-dhcp)# no disable`
Supported Releases	10.2.0E or later

dns-server address

Assigns a DNS server to clients based on the address pool.

Syntax	`dns-server address [address2...address8]`
Parameters	• *address* — Enter the DNS server IP address that services clients on the subnet in A.B.C.D or A::B format. • *address2...address8* — (Optional) Enter up to eight DNS server addresses, in order of preference.
Default	Not configured
Command Mode	DHCP-POOL
Usage Information	None
Example	`OS10(conf-dhcp-Dell)# dns-server 192.168.1.1`
Supported Releases	10.2.0E or later

domain-name

Configures the name of the domain where the device is located.

Syntax	`domain-name domain-name`
Parameters	*domain-name* — Enter the name of the domain (up to 32 characters).
Default	Not configured
Command Mode	DHCP-POOL
Usage Information	This is the default domain name that appends to hostnames that are not fully qualified. The `no` version of this command removes the configuration.
Example	`OS10(conf-dhcp-Dell)# domain-name dell.com`
Supported Releases	10.2.0E or later

hardware-address

Configures the client hardware address for manual configurations.

Syntax	`hardware-address nn:nn:nn:nn:nn:nn`
Parameters	*nn:nn:nn:nn:nn:nn* — Enter the 48-bit hardware address.
Default	Not configured
Command Mode	DHCP-POOL
Usage Information	The client hardware address is the MAC address of the client machine to which to lease a static IP address from.
Example	`OS10(conf-dhcp-static)# hardware-address 00:01:e8:8c:4d:0a`
Supported Releases	10.2.0E or later

host

Assigns a host to a single IPv4 or IPv6 address pool for manual configurations.

Syntax	`host A.B.C.D/A::B`
Parameters	*A.B.C.D/A::B* — Enter the host IP address in A.B.C.D or A::B format.
Default	Not configured
Command Mode	DHCP-POOL
Usage Information	The host address is the IP address used by the client machine for DHCP.
Example	`OS10(conf-dhcp-Dell)# host 20.1.1.100`
Supported Releases	10.2.0E or later

ip dhcp server

Enters DHCP mode.

Syntax	`ip dhcp server`
Parameters	None
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	This command is used to enter DHCP mode.
Example	`OS10(config)# ip dhcp server` `OS10(conf-dhcp)#`
Supported Releases	10.2.0E or later

ip helper-address

Configure the DHCP server address. Forwards UDP broadcasts received on an interface to the DHCP server. You can configure multiple helper addresses per interface by repeating the same command for each DHCP server address.

Syntax	`ip helper-address address [vrf vrf-name]`
Parameters	• *address* — Enter the IPv4 address to forward UDP broadcasts to the DHCP server in A.B.C.D format. • `vrf` *vrf-name* — (Optional) Enter `vrf` and then the name of the VRF through which the host address is reached.
Default	Disabled
Command Mode	INTERFACE
Usage Information	The DHCP server is available on L3 interfaces only. The `no` version of this command returns the value to the default.The client-facing and server-facing interfaces must be in the same VRF.
Example (IPv4)	`OS10(config)# interface eth 1/1/22` `OS10(conf-if-eth1/1/22)# ip helper-address 20.1.1.1 vrf blue`

Supported Releases 10.2.0E or later

ipv6 helper-address

Configure the DHCPv6 server address. Forwards UDP broadcasts received from IPv6 clients to the DHCPv6 server. You can configure multiple helper addresses per interface by repeating the same command for each DHCPv6 server address.

Syntax `ipv6 helper-address ipv6-address [vrf vrf-name]`

Parameters

- `vrf vrf-name` — (Optional) Enter the keyword `vrf` and then the name of the VRF through which the host address can be reached.
- *ipv6-address* — Specify the DHCPv6 server address in the A::B format.

Defaults Disabled

Command Mode INTERFACE

Usage Information The `no` version of this command deletes the IPv6 helper address.

Use this command on the interfaces where the DHCPv6 clients connect to forward the packets from clients to DHCPv6 server and vice-versa.

Example

```
OS10(config)# interface ethernet 1/1/22
OS10(conf-if-eth1/1/22)# ipv6 helper-address 2001:db8:0:1:1:1:1:1 vrf blue
```

Supported Releases 10.4.1.0 or later

lease

Configures a lease time for the IP addresses in a pool.

Syntax `lease {infinite | days [hours] [minutes]}`

Parameters

- `infinite` — Enter the keyword to configure a lease which never expires.
- *days* — Enter the number of lease days (0 to 31).
- *hours* — Enter the number of lease hours (0 to 23).
- *minutes* — Enter the number of lease minutes (0 to 59).

Default 24 hours

Command Mode DHCP-POOL

Usage Information The `no` version of this command removes the lease configuration.

Example

```
OS10(conf-dhcp-Dell)# lease 2 5 10
```

Example (Infinite)

```
OS10(conf-dhcp-Dell)# lease infinite
```

Supported Releases 10.2.0E or later

netbios-name-server address

Configures a NetBIOS WINS server which is available to DHCP clients.

Syntax `netbios-name-server ip-address [address2...address8]`

Parameters *ip-address* — Enter the address of the NetBIOS WINS server.

address2...address8 — (Optional) Enter additional server addresses.

Default Not configured

Command Mode DHCP-POOL

Usage Information Configure up to eight NetBIOS WINS servers available to a Microsoft DHCP client, in order of preference. The `no` version of this command returns the value to the default.

Example

```
OS10(conf-dhcp-Dell)# netbios-name-server 192.168.10.5
```

Supported Releases 10.2.0E or later

netbios-node-type

Configures the NetBIOS node type for the DHCP client.

Syntax `netbios-node-type type`

Parameters *type* — Enter the NetBIOS node type:

- `Broadcast` — Enter `b-node`.
- `Hybrid` — Enter `h-node`.
- `Mixed` — Enter `m-node`.
- `Peer-to-peer` — Enter `p-node`.

Default Hybrid

Command Mode DHCP-POOL

Usage Information The `no` version of this command resets the value to the default.

Example

```
OS10(conf-dhcp-Dell)# netbios-node-type h-node
```

Supported Releases 10.2.0E or later

network

Configures a range of IPv4 or IPv6 addresses in the address pool.

Syntax `network address/mask`

Parameters *address/mask* — Enter a range of IP addresses and subnet mask in A.B.C.D/x or A::B/x format.

Default Not configured

Command Mode DHCP-POOL

Usage Information	Use this command to configure a range of IPv4 or IPv6 addresses.
Example	`OS10(config-dhcp-Dell)# network 20.1.1.1/24`
Supported Releases	10.2.0E or later

pool

Creates an IP address pool name.

Syntax	`pool pool-name`
Parameters	*pool-name* — Enter the DHCP server pool name.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	Use this command to create an IP address pool name.
Example	`OS10(conf-dhcp)# pool Dell` `OS10(conf-dhcp-Dell)#`
Supported Releases	10.2.0E or later

range

Configures a range of IP addresses.

Syntax	`range {ip-address1 [ip-address2]}`
Parameters	• *ip-address1*—First IP address of the IP address range. • *ip-address2*—Last IP address of the IP address range.
Default	Not configured
Command Mode	DHCP-POOL
Usage Information	This command is used to configure a range of IP addresses that the OS10 switch, acting as the DHCP server, can assign to DHCP clients. The **no** version of this command requires only the first IP address to remove the range configuration.

Example	`OS10(config)# OS10(config)# ip dhcp server`
	`OS10(config-dhcp)# pool pool1`
	`OS10(config-dhcp-pool1)# network 192.168.10.0/24`
	`OS10(config-dhcp-pool1)# range 192.168.10.2 192.168.10.8`

Supported Releases	10.4.1 or later

show ip dhcp binding

Displays the DHCP binding table with IPv4 addresses.

Syntax	`show ip dhcp binding`
Parameters	None

Default	Not configured
Command Mode	EXEC
Usage Information	Use this command to view the DHCP binding table.

Example	`OS10# show ip dhcp binding`
	` IP Address  Hardware address  Lease expiration Hostname`
	`+-------------------------------------------------`
	`11.1.1.254   00:00:12:12:12:12 Jan 27 2016 06:23:45`
	
	`Total Number of Entries in the Table = 1`

Supported Releases	10.2.0E or later

DNS commands

OS10 supports the configuration of a DNS host and domain parameters.

ip domain-list

Adds a domain name to the DNS list. This domain name appends to incomplete hostnames in DNS requests.

Syntax	`ip domain-list [vrf vrf-name] [server-name] name`
Parameters	• `vrf vrf-name` — (Optional) Enter `vrf` and then the name of the VRF to add a domain name to the DNS list corresponding to that VRF. • *`server-name`* — (Optional) Enter the server name to add a domain name to the DNS list. • *`name`* — Enter the name of the domain to append to the DNS list.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	There is a maximum of six domain names to the DNS list. Use this domain name to complete unqualified host names. The `no` version of this command removes a domain name from the DNS list.
Example	`OS10(config)# ip domain-list jay dell.com`
Supported Releases	10.2.0E or later

ip domain-name

Configures the default domain and appends to incomplete DNS requests.

Syntax	`ip domain-name [vrf vrf-name] server-name`
Parameters	• `vrf vrf-name` — (Optional) Enter `vrf` and then the name of the VRF to configure the domain corresponding to that VRF. • *`server-name`* — (Optional) Enter the server name the default domain uses.
Default	Not configured
Command Mode	CONFIGURATION

Usage Information	This domain appends to incomplete DNS requests. The `no` version of this command returns the value to the default.
Example	`OS10(config)# ip domain-name jay dell.com`
Supported Releases	10.2.0E or later

ip host

Configures mapping between the host name server and the IP address.

Syntax	`ip host [vrf vrf-name] [host-name] address`
Parameters	• `vrf vrf-name` — (Optional) Enter `vrf` and then the name of the VRF to configure the name server to IP address mapping for that VRF. • *`host-name`* — (Optional) Enter the name of the host. • *`address`* — Enter an IPv4 or IPv6 address of the name server in A.B.C.D or A::B format.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The name-to-IP address table uses this mapping information to resolve host names. The `no` version of this command disables the mapping.
Example	`OS10(config)# ip host dell 1.1.1.1`
Supported Releases	10.2.0E or later

ip name-server

Configures up to a three IPv4 or IPv6 addresses used for network name servers.

Syntax	`ip name-server ip-address [ip-address2 ip-address3]`
Parameters	• *`ip-address`* — Enter the IPv4 or IPv6 address of a domain name server to use for completing unqualified names (incomplete domain names that cannot be resolved). • *`ip-address2 ip-address3`* — (Optional) Enter up two additional IPv4 or IPv6 name servers, separated with a space.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	OS10 does not support sending DNS queries over a VLAN. DNS queries are sent out on all other interfaces, including the Management port. You can separately configure both IPv4 and IPv6 domain name servers. In a dual stack setup, the system sends both A (request for IPv4) and AAAA (request for IPv6) record requests to a DNS server even if you only configure this command. The `no` version of this command removes the IP name-server configuration.
Example	`OS10(config)# ip name-server 10.1.1.5`
Supported Releases	10.2.0E or later

show hosts

Displays the host table and DNS configuration.

Syntax	`show hosts [vrf vrf-name]`
Parameters	`vrf vrf-name` — Enter `vrf` then the name of the VRF to display DNS host information corresponding to that VRF.
Default	Not configured
Command Mode	EXEC
Usage Information	None
Example	

```
OS10# show hosts
Default Domain Name : dell.com
Domain List : abc.com
Name Servers : 1.1.1.1 20::2
=========================================
       Static Host to IP mapping Table
=========================================
Host                              IP-Address
-----------------------------------------
dell-pc1                            20.1.1.1
```

Supported Releases	10.2.0E or later

IPv4 DHCP limitations

This section lists the DHCP limitations.

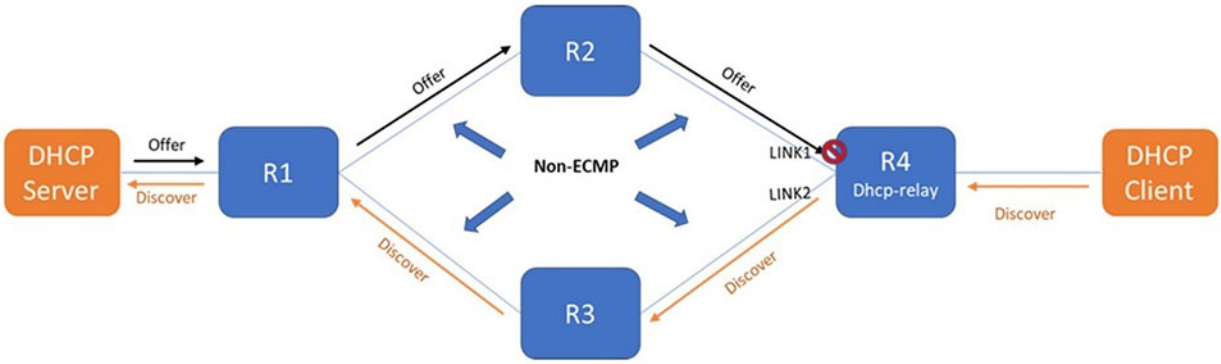
IPv4 DHCP asymmetric routing

OS10 does not support DHCP relay with IPv4 asymmetric routing. OS10 supports DHCP relay with IPv6 asymmetric routing.

The DHCP relay agent listens on the best DHCP server path. If the `DHCP OFFER` packet is sent through a path that the DHCP relay agent is not listening on, the `DHCP OFFER` packet drops. In this case, the DHCP client-enabled interface does not receive the IPv4 address.

For example, in the following topology, the `DHCP DISCOVER` packet is sent by the relay agent (R4) on link 2. The relay agent routing table points only to link 2 to reach the DHCP server. However, the DHCP server sends the `DHCP OFFER` packet to relay agent (R4) on a different path and the `DHCP OFFER` packet drops.

This issue occurs because the relay agent listens only on the *best path* uplink interfaces where the DHCP server is reachable.



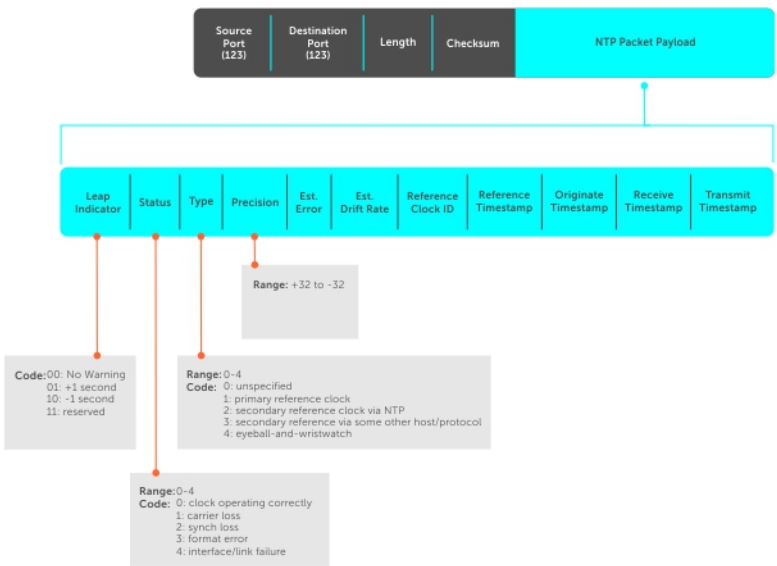
Network Time Protocol

Network Time Protocol (NTP) synchronizes timekeeping among a set of distributed time servers and clients. The protocol coordinates time distribution in a large, diverse network. NTP clients synchronize with NTP servers that provide accurate time measurement. NTP clients choose from several NTP servers to determine which offers the best available source of time and the most reliable transmission of information.

To get the correct time, OS10 synchronizes with a time-serving host. For the current time, you can set the system to poll specific NTP time-serving hosts. From those time-serving hosts, the system chooses one NTP host to synchronize with and acts as a client to the NTP host. After the host-client relationship establishes, the networking device propagates the time information throughout its local network.

The NTP client sends messages to one or more servers and processes the replies as received. Information included in the NTP message allows each client/server peer to determine the timekeeping characteristics of its other peers, including the expected accuracies of their clocks. Using this information, each peer selects the best time from several other clocks, updates the local clock, and estimates its accuracy.

NTP Fields



NOTE: OS10 supports both NTP server and client roles.

Enable NTP

NTP is disabled by default. To enable NTP, configure an NTP server where the system synchronizes. To configure multiple servers, enter the command multiple times. Multiple servers may impact CPU resources.

- Enter the IP address of the NTP server where the system synchronizes in CONFIGURATION mode.

  ```
  ntp server ip-address
  ```

View system clock state

```
OS10(config)# do show ntp status
system peer:          0.0.0.0
system peer mode:     unspec
leap indicator:       11
stratum:              16
precision:            -22
root distance:        0.00000 s
root dispersion:      1.28647 s
reference ID:         [73.78.73.84]
reference time:       00000000.00000000  Mon, Jan  1 1900  0:00:00.000
system flags:         monitor ntp kernel stats
jitter:               0.000000 s
stability:            0.000 ppm
broadcastdelay:       0.000000 s
authdelay:            0.000000 s
```

View calculated NTP synchronization variables

```
OS10(config)# do show ntp associations
     remote           local    st poll reach  delay   offset    disp
==============================================================================
 10.16.150.185   10.16.151.123 16 1024    0 0.00000  0.000000 3.99217
OS10# show ntp associations
     remote           local    st poll reach  delay   offset    disp
==============================================================================
 10.16.150.185   10.16.151.123 16 1024    0 0.00000  0.000000 3.99217
```

Broadcasts

Receive broadcasts of time information and set interfaces within the system to receive NTP information through broadcast. NTP is enabled on all active interfaces by default. If you disable NTP on an interface, the system drops any NTP packets sent to that interface.

1 Set the interface to receive NTP packets in INTERFACE mode.

  ```
  ntp broadcast client
  ```

2 Disable NTP on the interface in INTERFACE mode.

  ```
  ntp disable
  ```

Configure NTP broadcasts

```
OS10(config)# interface ethernet 1/1/10
OS10(conf-if-eth1/1/10)# ntp broadcast client
```

Disable NTP broadcasts

```
OS10(config)# interface ethernet 1/1/10
OS10(conf-if-eth1/1/10)# ntp disable
```

Source IP address

Configure one interface IP address to include in all NTP packets. The source address of NTP packets is the interface IP address the system uses to reach the network by default.

- Configure a source IP address for NTP packets in CONFIGURATION mode.
  ```
  ntp source interface
  ```
  - `ethernet` — Enter the keyword and node/slot/port information.
  - `port-channel` — Enter the keyword and number.
  - `vlan` — Enter the keyword and VLAN number, from 1 to 4093.
  - `loopback` — Enter the keyword and number, from 0 to 16383.
  - `mgmt` — Enter the keyword and node/slot/port information. The default is 1/1/1.

Configure source IP address

```
OS10(config)# ntp source ethernet 1/1/10
```

View source IP configuration

```
OS10(config)# do show running-configuration | grep source
ntp source ethernet1/1/1
```

Authentication

NTP authentication and the corresponding trusted key provides a reliable exchange of NTP packets with trusted time sources. NTP authentication begins with creating the first NTP packet after the key configuration. NTP authentication uses the message digest 5 (MD5) algorithm. The key is embedded in the synchronization packet that is sent to an NTP time source.

1. Enable NTP authentication in CONFIGURATION mode.
   ```
   ntp authenticate
   ```
2. Set an authentication key number and key in CONFIGURATION mode, from 1 to 4294967295.
   ```
   ntp authentication-key number md5 key
   ```
   - The *number* must match in the `ntp trusted-key` command.
   - The *key* is an encrypted string.
3. Define a trusted key in CONFIGURATION mode, from 1 to 4294967295. This *number* must match the number in the `ntp trusted-key` command.
   ```
   ntp trusted-key number
   ```
4. Configure an NTP server in CONFIGURATION mode.
   ```
   ntp server {hostname | ipv4-address | ipv6-address} [key keyid] [prefer]
   ```
   - *hostname* — Enter the keyword to see the IP address or host name of the remote device.
   - *ipv4-address* — Enter an IPv4 address in A.B.C.D format.
   - *ipv6-address* — Enter an IPv6 address in nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn format (elision of zeros is supported).
   - `key` *keyid* — Enter a text string as the key exchanged between the NTP server and the client.
   - `prefer` — Enter the keyword to set this NTP server as the preferred server.
5. Configure the NTP master and enter the stratum number that identifies the NTP server hierarchy in CONFIGURATION mode, from 2 to 10, default 8.
   ```
   ntp master <2-10>
   ```

Configure NTP

```
OS10(config)# ntp authenticate
OS10(config)# ntp trusted-key 345
OS10(config)# ntp authentication-key 345 mdf 0 5A60910FED211F02
OS10(config)# ntp server 1.1.1.1 key 345
OS10(config)# ntp master 7
```

View NTP configuration

```
OS10(config)# do show running-configuration
!
ntp authenticate
ntp authentication-key 345 mdf 0 5A60910FED211F02
ntp server 1.1.1.1 key 345
ntp trusted-key 345
ntp master 7
...
```

Sample NTP configuration

Following example shows an NTP master(11.0.0.2), server(10.0.0.1), and client(10.0.0.2) connected through a non-default VRF instance (VRF red). OS10 System acts as NTP server to synchronize its clock with the NTP master available in the non-default VRF instance red and provides time to NTP clients in that VRF instance.

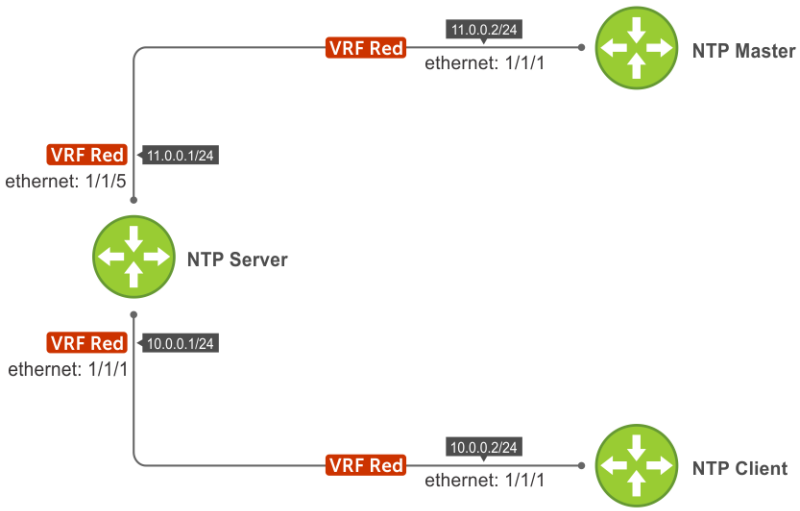


Figure 14. Sample NTP configuration

To create this sample NTP configuration, perform the following steps:

1 Configure the NTP server using the following steps:

a Create a non-default VRF instance and assign an interface to that VRF.

```
OS10(conf-vrf)# exit
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# ip vrf forwarding red
OS10(conf-if-eth1/1/1)# ip address 10.0.0.1/24
OS10(conf-if-eth1/1/1)# exit
OS10(config)# interface ethernet 1/1/5
OS10(conf-if-eth1/1/5)# no switchport
OS10(conf-if-eth1/1/5)# ip vrf forwarding red
```

```
OS10(conf-if-eth1/1/5)# ip address 11.0.0.1/24
OS10(conf-if-eth1/1/5)# exit
OS10(config)#
```

b Configure NTP master IP in NTP server

```
OS10(config)# ntp server 11.0.0.2
OS10(config)# do show running-configuration ntp
ntp server 11.0.0.2
OS10(config)#
```

NOTE: NTP master 11.0.0.2 is reachable only through red VRF.

c Configure NTP in red VRF instance.

```
OS10(config)# ntp enable vrf red

"% Warning: NTP server/client will be disabled in default VRF and enabled on a red VRF"
Do you wish to continue? (y/n): y

OS10(config)#
OS10(config)# do show running-configuration ntp
ntp server 11.0.0.2
ntp enable vrf red
OS10(config)#
```

2 Configure an NTP client using the following steps:

a Create non-default VRF instance and assign interface to the VRF.

```
OS10(config)# ip vrf red
OS10(conf-vrf)# exit
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# ip vrf forwarding red
OS10(conf-if-eth1/1/1)# ip address 10.0.0.2/24
OS10(conf-if-eth1/1/1)# exit
OS10(config)#
```

b Configure NTP server IP in NTP client.

```
OS10(config)# ntp server 10.0.0.1
OS10(config)# do show running-configuration ntp
ntp server 10.0.0.1
OS10(config)#
```

c Configure NTP in red VRF instance.

```
OS10(config)# ntp enable vrf red

"% Warning: NTP server/client will be disabled in default VRF and enabled on a red VRF"
Do you wish to continue? (y/n): y

OS10(config)# do show running-configuration ntp
ntp server 10.0.0.1
ntp enable vrf red
OS10(config)#
```

3 Configure an NTP master using the following steps:

a Create non-default VRF instance and assign interface to the VRF.

```
OS10(config)# ip vrf red
OS10(conf-vrf)# exit
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# ip vrf forwarding red
OS10(conf-if-eth1/1/1)# ip address 11.0.0.2/24
OS10(conf-if-eth1/1/1)# exit
OS10(config)
```

b Configure NTP as master.

```
OS10(config)# ntp master
OS10(config)# do show running-configuration ntp
ntp master 8
OS10(config)#
```

c Configure NTP in red VRF instance.

```
OS10(config)# ntp enable vrf red

"% Warning: NTP server/client will be disabled in default VRF and enabled on a red VRF"
Do you wish to continue? (y/n): y

OS10(config)# do show running-configuration ntp
ntp master 8
ntp enable vrf red
OS10(config)#
```

4 Verify NTP client(10.0.0.2) is connected to NTP server(10.0.0.1) running in red VRF.

```
OS10# show ntp associations vrf red

     remote           refid      st t when poll reach   delay   offset  jitter

==============================================================================

*10.0.0.1        11.0.0.2        10 u    2   64    1    0.578   -1.060   0.008

OS10# show ntp status vrf red
associd=0 status=0615 leap_none, sync_ntp, 1 event, clock_sync,
system peer:        10.0.0.1:123
system peer mode:   client
leap indicator:     00
stratum:            11
log2 precision:     -24
root delay:         0.991
root dispersion:    1015.099
reference ID:       10.0.0.1
reference time:     dbc7b087.5d47aaa6  Sat, Nov  5 2016  1:12:39.364
system jitter:      0.000000
clock jitter:       0.462
clock wander:       0.003
broadcast delay:    -50.000
symm. auth. delay:  0.000
OS10#
```

5 Verify NTP server (10.0.0.1) is connected to NTP master(11.0.0.2) running in red VRF.

```
OS10(config)# do show ntp associations vrf red

     remote           refid      st t when poll reach   delay   offset  jitter

==============================================================================

 LOCAL(0)        .LOCL.           8 l  111   64    2    0.000    0.000   0.000

*11.0.0.2        LOCAL(0)         9 u   43   64    3    0.441    0.026   0.047

OS10(config)# do show ntp status vrf red
associd=0 status=0615 leap_none, sync_ntp, 1 event, clock_sync,
system peer:        11.0.0.2:123
system peer mode:   client
leap indicator:     00
stratum:            10
log2 precision:     -24
root delay:         0.441
root dispersion:    950.580
reference ID:       11.0.0.2
reference time:     dbc7b03e.733f51d7  Sat, Nov  5 2016  1:11:26.450
system jitter:      0.000000
clock jitter:       0.009
clock wander:       0.000
broadcast delay:    -50.000
symm. auth. delay:  0.000
OS10(config)#
```

NTP commands

ntp authenticate

Enables authentication of NTP traffic between the device and the NTP time serving hosts.

Syntax	`ntp authenticate`
Parameters	None
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	You must also configure an authentication key for NTP traffic using the `ntp authentication-key` command. The `no` version of this command disables NTP authentication.
Example	`OS10(config)# ntp authenticate`
Supported Releases	10.2.0E or later

ntp authenticate-key

Configures the authentication key for trusted time sources.

Syntax	`ntp authenticate-key` *`number`* `md5` *`[0 | 7] key`*
Parameters	• *`number`* — Enter the authentication key number, from 1 to 4294967295. • `md5` — Set to MD5 encryption. • `0` — Set to unencrypted format, the default. • `7` — Set to hidden encryption. • *`key`* — Enter the authentication key.
Default	0
Command Mode	CONFIGURATION
Usage Information	The authentication number must be the same as the `number` parameter configured in the `ntp trusted-key` command. Use the `ntp authenticate` command to enable NTP authentication.
Example	`OS10(config)# ntp authentication-key 1200 md5 0 dell`
Supported Releases	10.2.0E or later

ntp broadcast client

Configures the interface to receive NTP broadcasts from an NTP server.

Syntax	`ntp broadcast client`
Parameters	None

Default	Not configured
Command Mode	INTERFACE
Usage Information	The `no` version of this command disables broadcast.
Example	`OS10(conf-if-eth1/1/1)# ntp broadcast client`
Supported Releases	10.2.0E or later

ntp disable

By default, NTP is enabled on all interfaces. Prevents an interface from receiving NTP packets.

Syntax	`ntp disable`
Parameters	None
Default	Enabled
Command Mode	INTERFACE
Usage Information	USe this command to configure OS10 to not listen to a particular server and prevent the interface from receiving NTP packets. The `no` version of this command re-enables NTP on an interface.
Example	`OS10(conf-if-eth1/1/7)# ntp disable`
Supported Releases	10.2.0E or later

ntp enable vrf

Enables NTP for the management or non-default VRF instance.

Syntax	`ntp enable vrf {management | vrf` *`vrf-name`*`}`
Parameters	• `management` — Enter the keyword to enable NTP for the management VRF instance. • `vrf` *`vrf-name`* — Enter the keyword then the name of the VRF to enable NTP for that non-default VRF instance.
Defaults	Disabled
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command disables NTP for the management VRF instance.
Example	`OS10(config)# ntp enable vrf management` `OS10(config)# ntp enable vrf vrf-blue`
Supported Releases	10.4.0E(R1) or later

ntp master

Configures an NTP master server.

Syntax	`ntp master` *`stratum`*
Parameters	*`stratum`* — Enter the stratum number to identify the NTP server hierarchy, from 2 to 10.

Default	8
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command resets the value to the default.
Example	`OS10(config)# ntp master 6`
Supported Releases	10.2.0E or later

ntp server

Configures an NTP time-serving host.

Syntax	`ntp server {hostname | ipv4-address | ipv6-address} [key keyid] [prefer]`
Parameters	• *hostname* — Enter the host name of the server. • *ipv4-address | ipv6-address* — Enter the IPv4 address in A.B.C.D format or IPv6 address in A::B format of the NTP server. • `key` *keyid* — (Optional) Enter the NTP peer key ID, from 1 to 4294967295. • `prefer` — (Optional) Configures this peer to have priority over other servers.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	You can configure multiple time-serving hosts. From these time-serving hosts, the system chooses one NTP host to synchronize with. To determine which server to select, use the `show ntp associations` command. Dell EMC recommends limiting the number of hosts you configure, as many polls to the NTP hosts can impact network performance.
Example	`OS10(config)# ntp server eureka.com`
Supported Releases	10.2.0E or later

ntp source

Configures an interface IP address to include in NTP packets.

Syntax	`ntp source interface`
Parameters	*interface* — Set the interface type: • `ethernet` *node/slot/port[:subport]* — Enter the Ethernet interface information. • `port-channel` *id-number* — Enter the port-channel number, from 1 to 128. • `vlan` *vlan-id* — Enter the VLAN number, from 1 to 4093. • `loopback` *loopback-id* — Enter the Loopback interface number, from 0 to 16383. • `mgmt` *node/slot/port* — Enter the Management port interface information.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command removes the configuration.

Example

```
OS10(config)# ntp source ethernet 1/1/24
```

Supported Releases 10.2.0E or later

ntp trusted-key

Sets a key to authenticate the system to which NTP synchronizes with.

Syntax `ntp trusted-key number`

Parameters *number* — Enter the trusted key ID (1 to 4294967295).

Default Not configured

Command Mode CONFIGURATION

Usage Information The `number` parameter must be the same number as the `number` parameter in the `ntp authentication-key` command. If you change the `ntp authentication-key` command, you must also change this command. The `no` version of this command removes the key.

Example

```
OS10(config)# ntp trusted-key 234567
```

Supported Releases 10.2.0E or later

show ntp associations

Displays the NTP master and peers.

Syntax `show ntp associations [vrf {management | vrf vrf-name}]`

Parameters

- `management` — Enter the keyword to display NTP information corresponding to the management VRF instance.
- `vrf vrf-name` — Enter the keyword then the name of the VRF to display NTP information corresponding to that non-default VRF instance.

Default Not configured

Command Mode EXEC

Usage Information

- `(none)` — One or more of the following symbols displays:
  - `*` — Synchronized to this peer.
  - `#` — Almost synchronized to this peer.
  - `+` — Peer was selected for possible synchronization.
  - `-` — Peer is a candidate for selection.
  - `~` — Peer is statically configured.
- `remote` — Remote IP address of the NTP peer.
- `ref clock` — IP address of the remote peer reference clock.
- `st` — Peer stratum, the number of hops away from the external time source. `16` means that the NTP peer cannot reach the time source.
- `when` — Last time the device received an NTP packet.
- `poll` — Polling interval in seconds.
- `reach` — Reachability to the peer in octal bitstream.

- `delay` — Time interval or delay for a packet to complete a round-trip to the NTP time source in milliseconds.
- `offset` — Relative time of the NTP peer's clock to the network device clock in milliseconds.
- `disp` — Dispersion.

Example

```
OS10# show ntp associations
remote      ref clock  st when poll reach delay  offset disp
==============================================================
 10.10.120.5 0.0.0.0     16 - 256        0 0.00 0.000 16000.0
*172.16.1.33 127.127.1.0 11 6 16         377   -0.08 -1499.9 104.16
 172.31.1.33 0.0.0.0     16 - 256        0 0.00 0.000 16000.0
 192.200.0.2 0.0.0.0     16 - 256        0 0.00 0.000 16000.0
```

```
OS10# show ntp associations vrf management
remote     local      st    poll reach  delay       offset        disp
==============================================================================
*1.1.1.2   1.1.1.1     3    64    1     0.00027     0.000056    0.43309
```

Supported Releases 10.2.0E or later

show ntp status

Displays NTP configuration information.

Syntax

```
show ntp status [vrf {management | vrf vrf-name]
```

Parameters

- *status* — (Optional) View the NTP status.
- `vrf management` — (Optional) Enter the keywords to display NTP information corresponding to the management VRF.
- `vrf vrf-name` — (Optional) Enter the keyword then the name of the VRF to display NTP information corresponding to that non-default VRF.

Default Not configured

Command Mode EXEC

Usage Information None

Example (Status)

```
OS10# show ntp status
system peer:       0.0.0.0
system peer mode:  unspec
leap indicator:    11
stratum:           16
precision:         -22
root distance:     0.00000 s
root dispersion:   1.28647 s
reference ID:      [73.78.73.84]
reference time:    00000000.00000000 Mon, Jan 1 1900 0:00:00.000
system flags:      monitor ntp kernel stats
jitter:            0.000000 s
stability:         0.000 ppm
broadcastdelay:    0.000000 s
authdelay:         0.000000 s
```

```
OS10# show ntp status vrf management
system peer:          1.1.1.2
system peer mode:     client
leap indicator:       00
stratum:              4
precision:            -23
root distance:        0.00027 s
root dispersion:      0.94948 s
reference ID:         [1.1.1.2]
```

```
reference time:       ddc78084.f17ea38b  Tue, Nov 28 2017  6:28:20.943
system flags:         ntp kernel stats
jitter:               0.000000 s
stability:            0.000 ppm
broadcastdelay:       0.000000 s
authdelay:            0.000000 s
OS10#
```

```
OS10# show ntp status vrf red
associd=0 status=0618 leap_none, sync_ntp, 1 event, no_sys_peer,
system peer:        11.0.0.2:123
system peer mode:   client
leap indicator:     00
stratum:            10
log2 precision:     -24
root delay:         0.338
root dispersion:    1136.790
reference ID:       11.0.0.2
reference time:     dbc7a951.f7978096  Sat, Nov  5 2016  0:41:53.967
system jitter:      0.000000
clock jitter:       0.003
clock wander:       0.001
broadcast delay:    -50.000
symm. auth. delay:  0.000
```

Supported Releases 10.2.0E or later

System clock

OS10 uses the network time protocol (NTP) to synchronize the system clock with a time-serving host. If you do not use NTP, set the system time and the timezone. The hardware-based real-clock time (RTC) is reset to the new system time.

You can set the current time and date after you disable NTP. When you enable NTP, it overwrites the system time.

- Enter the time and date in EXEC mode.
  ```
  clock set time year-month-day
  ```
  Enter *time* in the format *hour*:*minute*:*second*, where *hour* is 1 to 24; *minute* is 1 to 60; *second* is 1 to 60. For example, enter 5:15 PM as `17:15:00`.

  Enter *year-month-day* in the format `YYYY-MM-DD`, where `YYYY` is a four-digit year, such as `2016`; `MM` is a month from 1 to 12; `DD` is a day from 1 to 31.
- Enter the timezone in CONFIGURATION mode.
  ```
  clock timezone timezone-string Hours Minutes
  ```
  Enter *timezone-string* which is the name of the time zone.

  Enter *Hours* offset from UTC, ranging from **–23** to **23**.

  Enter *Minutes* offset from UTC, ranging from **0** to **59**.

Set time and date

```
OS10# clock set 13:00:00 2018-08-30
```

View system time and date

```
OS10# show clock
2018-08-30T13:01:01.45+00:00
```

Set time zone

```
OS10(config)# clock timezone IST 5 30
```

View system time and date with time zone configured

```
OS10# show clock
2018-08-30T13:01:01.57+05:30
```

System Clock commands

clock set

Sets the system time.

Syntax `clock set time year-month-day`

Parameters

time	Enter *time* in the format *hour*:*minute*:*second*, where *hour* is 1 to 24; *minute* is 1 to 60; *second* is 1 to 60. For example, enter 5:15 PM as `17:15:00`.
year-month-day	Enter *year-month-day* in the format `YYYY-MM-DD`, where `YYYY` is a four-digit year, such as `2016`; `MM` is a month from 1 to 12; `DD` is a day from 1 to 31.

Default Not configured

Command Mode EXEC

Usage Information Use this command to reset the system time if the system clock is out of synch with the NTP time. The hardware-based real-clock time (RTC) resets to the new time. The new system clock setting applies immediately.

Example

```
OS10# clock set 18:30:10 2017-01-25
```

Supported Releases 10.2.1E or later

clock timezone

Sets the time zone used for the system clock.

Syntax `clock timezone timezone-string Hours Minutes`

Parameters

- Enter *timezone-string* which is the name of the time zone.
- Enter *Hours* offset from UTC, ranging from **–23** to **23**.
- Enter *Minutes* offset from UTC, ranging from **0** to **59**.

Default Not configured

Command Mode CONFIGURATION

Usage Information Universal time coordinated (UTC) is the time standard based on Greenwich Mean time. To set the time zone for the system clock, enter the difference of hours between UTC and your time zone.

Example

```
OS10(config)# clock timezone IST 5 30
```

Supported Releases 10.3.0E or later

show clock

Displays the current system clock settings.

Syntax	`show clock`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	The universal time coordinated (UTC) value is the number of hours that your time zone is later or earlier than UTC/ Greenwich mean time.
Example	`OS10# show clock` `2017-01-25T11:00:31.68-08:00`
Supported Releases	10.2.1E or later

System banners

You can configure a system login and message of the day (MOTD) text banners. The system login banner displays before you log in. The MOTD banner displays immediately after a successful login.

You can reset the banner text to the Dell EMC default banner or disable the banner display.

Login banner

Configure a system login banner that displays before you log in using interactive mode. Starting and ending double-quotes are not necessary. Enter a single delimiter character or the key combination ^C to specify the start and end of the text banner.

1 Enter the `banner login` command with a single delimiter character and press **Enter**.
2 Enter each line of text and press **Enter**. Enter a maximum of 4096 characters. There is no limit to the number of lines.
3 Complete the banner configuration by entering a line that contains only the delimiter character.

- Enable a login banner in CONFIGURATION mode. Enclose the delimiters and banner text in double-quotes.

```
banner login delimiter <Enter>
banner-text <Enter>
banner-text <Enter>
... <Enter>
delimiter
```

Configure login banner

```
OS10(config)# banner login %
DellEMC S4148U-ON login
Enter your username and password
%
```

To delete a login banner and reset it to the Dell EMC default banner, enter the `no banner login` command. To disable banner display before login, enter the `banner login disable` command.

MOTD banner

Configure a message of the day banner that displays after you log in. Enter up to 4096 characters. To start and end the MOTD banner, enter a single delimiter character or the key combination ^C. You can enter any character as the delimiter.

To enter a MOTD banner text, use the interactive mode. Enter the command with the delimiter character and press **Enter**. Then enter each line and press **Enter**. Complete the banner configuration by entering a line that contains only the delimiter character. Starting and ending double-quotes are not necessary.

Configure MOTD banner

```
OS10(config)# banner motd %
DellEMC S4148U-ON
Today's tip: Press tab or spacebar for command completion.
Have a nice day!
%
```

To delete a MOTD banner and reset it to the Dell EMC default MOTD banner, enter the `no banner motd` command. To disable MOTD banner display after login, enter the `banner motd disable` command.

System banner commands

banner login

Configures a login banner that displays before you log in to the system.

Syntax

```
banner login delimiter <Enter>
banner-text <Enter>
banner-text <Enter>
... <Enter>
delimiter
```

Parameters

- *delimiter* — Enter a single delimiter character or the key combination ^C to specify the start and end of the text banner.
- *banner-text* — Enter a maximum of 4096 characters. There is no limit to the number of lines.

Default The Dell EMC default banner is displayed before you log in.

Command Mode CONFIGURATION

Usage Information

- To enter multiline banner text, use the interactive mode. Enter the command with the delimiter character and press **Enter**. Then enter each line and press **Enter**. Complete the banner configuration by entering a line that contains only the delimiter character. Starting and ending double-quotes are not necessary.
- To delete a login banner and reset it to the Dell EMC default banner, enter the `no banner login` command. To disable banner display before login, enter the `banner login disable` command.

Example

```
OS10(config)# banner login %
Welcome to DellEMC Z9100-ON
Enter your username and password
%
```

Supported Releases 10.4.1.0 or later

banner motd

Configures a multi-line message of the day banner that displays after you log in.

Syntax

```
banner motd delimiter <Enter>
banner-text <Enter>
banner-text <Enter>
... <Enter>
delimiter
```

Parameters

- *delimiter* — Enter a single delimiter character or the key combination ^C to specify the start and end of the text banner.
- *banner-text* — Enter a maximum of 4096 characters. There is no limit on the number of lines.

Default The Dell EMC default MOTD banner is displayed after you log in.

Command Mode CONFIGURATION

Usage Information

- To enter a MOTD banner text, use the interactive mode. Enter the command with the delimiter character and press **Enter**. Then enter each line and press **Enter**. Complete the banner configuration by entering a line that contains only the delimiter character. Starting and ending double-quotes are not necessary.
- To delete a login banner and reset it to the Dell EMC default banner, enter the `no banner motd` command. To disable banner display before login, enter the `banner motd disable` command.

Example

```
OS10(config)# banner motd %
DellEMC S4148U-ON
Today's tip: Press tab or spacebar for command completion.
Have a nice day!
%
```

Supported releases 10.4.1.0 or later

User session management

You can manage the active user sessions using the following commands:

- Configure the timeout for all the active user sessions using the `exec-timeout` *timeout-value* command in the CONFIGURATION mode.
- Clear any user session using the `kill-session` *session-ID* command in the EXEC mode.
- View the active user sessions using the `show sessions` command in the EXEC mode.

Configure timeout for user sessions

```
OS10(config)# exec-timeout 300
OS10(config)#
```

Clear user session

```
OS10# kill-session 3
```

View active user sessions

```
OS10# show sessions

Current session's operation mode: Non-transaction
```

```
Session-ID User      In-rpcs In-bad-rpcs Out-rpc-err Out-notify Login-time             Lock
-------------------------------------------------------------------------------------------
 3          snmp_user 114     0          0           0          2017-07-10T23:58:39Z
 4          snmp_user 57      0          0           0          2017-07-10T23:58:40Z
 6          admin     17      0          0           4          2017-07-12T03:55:18Z
*7          admin     10      0          0           0          2017-07-12T04:42:55Z
OS10#
```

User session management commands

exec-timeout

Configure timeout in seconds for all the user sessions.

Syntax	`exec-timeout` *timeout-value*
Parameters	*timeout-value* — Enter the timeout value in seconds, from 0 to 3600.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command disables the timeout.
Example	`OS10(config)# exec-timeout 300` `OS10(config)#`
Supported Releases	10.3.1E or later

kill-session

Terminate a user session.

Syntax	`kill-session` *session-ID*
Parameters	*session-ID* — Enter the user session ID.
Default	Not configured
Command Mode	EXEC
Usage Information	None
Example	`OS10# kill-session 3`
Supported Releases	10.3.1E or later

show sessions

Displays the active management sessions.

Syntax	`show sessions`
Parameters	None
Default	Not configured

Command Mode	EXEC
Usage Information	Use this command to view information about the active user management sessions.
Example	

```
OS10# show sessions

Current session's operation mode: Non-transaction

Session-ID User      In-rpcs In-bad-rpcs Out-rpc-err Out-notify Login-time          Lock
-------------------------------------------------------------------------------------------
 3         snmp_user 114     0           0           0          2017-07-10T23:58:39Z
 4         snmp_user 57      0           0           0          2017-07-10T23:58:40Z
 6         admin     17      0           0           4          2017-07-12T03:55:18Z
*7         admin     10      0           0           0          2017-07-12T04:42:55Z
OS10#
```

Supported Releases	10.3.1E or later

Telnet server

To allow Telnet TCP/IP connections to an OS10 switch, enable the Telnet server. The OS10 Telnet server uses the Debian `telnetd` package. By default, the Telnet server is disabled.

When you enable the Telnet server, connect to the switch using the IP address configured on the management or any front-panel port. The Telnet server configuration is persistent and is maintained after you reload the switch. To verify the Telnet server configuration, enter the `show running-configuration` command.

Enable Telnet server

```
OS10(config)# ip telnet server enable
```

Disable Telnet server

```
OS10(config)# no ip telnet server enable
```

By default, the Telnet server is disabled on the default virtual routing and forwarding (VRF). To configure the Telnet server to be reachable on the management VRF, use the `ip telnet server vrf management` command.

Configure Telnet server on management VRF

```
OS10(config)# ip telnet server vrf management
```

Telnet commands

ip telnet server enable

Enables Telnet TCP/IP connections to an OS10 switch.

Syntax	`ip telnet server enable`
Parameters	None
Default	Disabled
Command Mode	CONFIGURATION
Usage Information	By default, the Telnet server is disabled. When you enable the Telnet server, use the IP address configured on the management or any front-panel port to connect to an OS10 switch. After you reload the switch, the Telnet server

configuration is maintained. To verify the Telnet server configuration, use the `show running-configuration` command.

Example

```
OS10(config)# ip telnet server enable
```

Example (disable)

```
OS10(config)# no ip telnet server enable
```

Supported Releases 10.4.0E(R1) or later

ip telnet server vrf

Configures the Telnet server for the management or non-default VRF instance.

Syntax `ip telnet server vrf {management | vrf vrf-name}`

Parameters

- `management` — Configures the management VRF used to reach the Telnet server.
- `vrf vrf-name` — Enter the keyword `vrf` followed by the name of the VRF to configure the non-default VRF instance used to reach the Telnet server.

Default The Telnet server is reachable on the default VRF.

Command Mode CONFIGURATION

Usage Information By default, the Telnet server is disabled. To enable the Telnet server, use the `telnet enable` command.

Example

```
OS10(config)# ip telnet server vrf management
OS10(config)# ip telnet server vrf vrf-blue
```

Supported Releases 10.4.0E(R1) or later

Security

Authentication, authorization, and accounting (AAA) services secure networks against unauthorized access. In addition to local authentication, OS10 supports remote authentication dial-in user service (RADIUS) and terminal access controller access control system (TACACS+) client/server authentication systems. For RADIUS and TACACS+, an OS10 switch acts as a client and sends authentication requests to a server that contains all user authentication and network service access information.

A RADIUS or TACACS+ server provides authentication (user credentials verification), authorization (role-based permissions), and accounting services. You can configure the security protocol used for different login methods and users. RADIUS provides very limited authorization and accounting services compared to TACACS+.

AAA authentication methods

An OS10 switch uses a list of authentication methods to define the types of authentication and the sequence in which they apply. By default, only the `local` authentication method is configured.

The authentication methods in the method list execute in the order in which you configure them. You can re-enter the methods to change the order. The `local` authentication method remains enabled even if you remove all configured methods in the list using the `no aaa authentication login {console | default}` command.

- Configure the AAA authentication method in CONFIGURATION mode.
  `aaa authentication login {console | default} {local | group radius | group tacacs+}`
  - `console` — Configure authentication methods for console logins.
  - `default` — Configure authentication methods for non-console such as SSH and Telnet logins.

- `local` — Use the local username, password, and role entries configured with the `username password role` command.
- `group radius` — Use the RADIUS servers configured with the `radius-server host` command.
- `group tacacs+` — Use the TACACS+ servers configured with the `tacacs-server host` command.

Configure user role on server

If a console user logs in with RADIUS or TACACS+ authentication, the role you configured for the user on the RADIUS or TACACS+ server applies. User authentication fails if no role is configured on the authentication server.

In addition, you must configure the user role on the RADIUS or TACACS+ server using the vendor-specific attribute (VSA) or the authentication fails. Dell's vendor ID is 674. You create a VSA with `Name = Dell-group-name, OID = 2, Type = string`. Valid values for *Dell-group-name* are `sysadmin`, `secadmin`, `netadmin`, and `netoperator`. Use the VSA *Dell-group-name* values when you create users on a Radius or TACACS+ server.

For detailed information about how to configure vendor-specific attributes on a RADIUS or TACACS+ server, refer to the respective RADIUS or TACACS+ server documentation.

Configure AAA authentication

```
OS10(config)# aaa authentication login default group radius local
OS10(config)# do show running-configuration aaa
aaa authentication login default group radius local
aaa authentication login console local
```

Remove AAA authentication methods

```
OS10(config)# no aaa authentication login default
OS10(config)# do show running-configuration aaa
aaa authentication login default local
aaa authentication login console local
```

User re-authentication

To prevent users from accessing resources and performing tasks for which they are not authorized, OS10 allows you to require users to re-authenticate by logging in again when an authentication method or server changes, such as:

- Adding or removing a RADIUS server using the `radius-server host` command
- Adding or removing an authentication method using the `aaa authentication login {console | default} {local | group radius | group tacacs+}` command

You can enable this feature so that user re-authentication is required when any of these actions are performed. In these cases, logged-in users are logged out of the switch and all OS10 sessions terminate. By default, user re-authentication is disabled.

Enable user re-authentication

- Enable user re-authentication in CONFIGURATION mode.

  ```
  aaa re-authenticate enable
  ```

The `no` version of this command disables user re-authentication.

Password strength

By default, the password you configure with the `username password` command must be at least nine alphanumeric characters.

To increase password strength, you can create password rules using the `password-attributes` command. When you enter the command, at least one parameter is required. When you enter the `character-restriction` parameter, at least one option is required.

- Create rules for stronger passwords in CONFIGURATION mode.

```
password-attributes {[min-length number] [character-restriction {[upper number]
[lower number][numeric number] [special-char number]}}
```

  - `min-length number` — Enter the minimum number of required alphanumeric characters, from 6 to 32; default 9.
  - `character-restriction` — Enter a requirement for the alphanumeric characters in a password:
    - `upper number` — Minimum number of uppercase characters required, from 0 to 31; default 0.
    - `lower number` — Minimum number of lowercase characters required, from 0 to 31; default 0.
    - `numeric number` — Minimum number of numeric characters required, from 0 to 31; default 0.
    - `special-char number` — Minimum number of special characters required, from 0 to 31; default 0.

Create password rules

```
OS10(config)# password-attributes min-length 7 character-restriction upper 4 numeric 2
```

Display password rules

```
OS10(config)# do show running-configuration password-attributes
password-attributes min-length 7 character-restriction upper 4 numeric 2
```

Role-based access control

RBAC provides control for access and authorization. Users are granted permissions based on defined roles — not on their individual system user ID. Create user roles based on job functions to help users perform their associated job function. You can assign each user only a single role, and many users can have the same role. A user role authenticates and authorizes a user at login, and places you in EXEC mode (see CLI basics).

OS10 supports four pre-defined roles: sysadmin, secadmin, netadmin, and netoperator. Each user role assigns permissions that determine the commands a user can enter, and the actions a user can perform. RBAC provides an easy and efficient way to administer user rights. If a user's role matches one of the allowed user roles for a command, command authorization is granted.

The OS10 RBAC model provides separation of duty as well as greater security. It places some limitations on each role's permissions to allow you to partition tasks. For greater security, only some user roles can view events, audits, and security system logs.

Assign user role

To limit OS10 system access, assign a role when you configure each user.

- Enter a user name, password, and role in CONFIGURATION mode.

```
username username password password role role
```

  - `username username` — Enter a text string. A maximum of 32 alphanumeric characters; 1 character minimum.
  - `password password` — Enter a text string. A maximum of 32 alphanumeric characters; 9 characters minimum.
  - `role role` — Enter a user role:
    - `sysadmin` — Full access to all commands in the system, exclusive access to commands that manipulate the file system, and access to the system shell. A system administrator can create user IDs and user roles.
    - `secadmin` — Full access to configuration commands that set security policy and system access, such as password strength, AAA authorization, and cryptographic keys. A security administrator can display security information, such as cryptographic keys, login statistics, and log information.
    - `netadmin` — Full access to configuration commands that manage traffic flowing through the switch, such as routes, interfaces, and ACLs. A network administrator cannot access configuration commands for security features or view security information.

- `netoperator` — Access to EXEC mode to view the current configuration. A network operator cannot modify any configuration setting on a switch.

Create user and assign role

```
OS10(config)# username smith password silver403! role sysadmin
```

View users

```
OS10# show users

Index Line  User    Role     Application Idle  Login-Time             Location
----- ----  ------  ------   ----------- ----  --------------------   -------------
1     ttyS  root    root     -bash       >24h  2018-05-23 T23:05:03Z  console
2     pts/0 admin   sysadmin bash        1.1s  2018-05-30 T20:04:27Z  10.14.1.214[ssh]
```

Bootloader Protection

Protecting the bootloader via a GRUB password is essential to prevent unauthorised users with malicious intent from accessing your switch. OS10 provides a set of three commands which allow you to enable, disable or view bootloader protection information.

This feature is available only for the `sysadmin` and `secadmin` roles.

WARNING: When you enable this feature ensure to keep a copy of a configured username and password, as you cannot recover the switch without the configured credentials.

- To enable bootloader protection, use the `boot protect enable username` *`username`* `password` *`password`* command. This command allows you to setup a username and password for bootloader protection. You can configure a maximum of three users per console.

  boot protect enable username password

  ```
  OS10# boot protect enable username root password calvin
  ```

- To disable bootloader protection, use the `boot protect disable username` command. This command allows you to disable bootloader protection by username.

  boot protect disable username

  ```
  OS10# boot protect disable username root
  ```

- To display information about the current list of users configured for bootloader protection, use the `show boot protect` command.

  show boot protect (when disabled)

  ```
  OS10# show boot protect
  Boot protection disabled
  ```

  show boot protect (when enabled)

  ```
  OS10# show boot protect
  Boot protection enabled
  Authorized users: root linuxadmin admin
  ```

Linuxadmin User Configuration

OS10 supports two factory default users, the `admin` and `linuxadmin`. The `admin` user is used to sign into the CLI and the `linuxadmin` user is used to access the linux shell.

Support has now been added to manage the default `linuxadmin` user via the CLI. You can perform the following operations via the CLI:

NOTE:

- This feature is used to manage the default linux user, `linuxadmin` only. This feature cannot be used to manage linux users created via the root user.
- The password of the `linuxadmin` user that is configured via the OS10 CLI takes higher precedence over the password modified via the linux shell across reboots.
- The password of the `linuxadmin` user configured via the linux shell does not show up in the `show running-configuration` output.

Configure a password for the linuxadmin user:

To configure a password for the `linuxadmin` user, enter CONFIGURATION mode and execute the command `system-user linuxadmin password {clear-text-password | hashed-password}`.

```
OS10(config)# system-user linuxadmin password Dell@Force10!@
OS10(config)# exit
OS10# write memory
OS10#
OS10# exit
```

```
OS10(config)# system-user linuxadmin password $6$3M55wOYy
$Sw1V9Ok3GE4Hmf6h1ARH.dBHy9gpEFYUvdu15ZpnCYzt.nJjFm0VIz/rQvvJeX6krRtfYs2ZqBl6TkmLGAwtM
OS10(config)# exit
OS10# write memory
OS10#
OS10# exit
```

Display the linuxadmin password:

To display the `linuxadmin` password, use the `show running-configuration` command.

```
OS10# show running-configuration | grep system-user
system-user linuxadmin password
$6$5DdOHYg5$JCE1vMSmkQOrbh31U74PIPv7lyOgRmba1IxhkYibppMXs1KM4Y.gbTPcxyMP/PHUkMc5rdk/
ZLv9Sfv3ALtB61
OS10#
```

Disabling or locking the linuxadmin user:

To disable or lock the `linuxadmin` user, enter CONFIGURATION mode and execute the command `system-user linuxadmin disable`.

```
OS10(config)# system-user linuxadmin disable
OS10(config)#
```

Enabling or unlocking the linuxadmin user:

To enable or unlock the `linuxadmin` user, enter CONFIGURATION mode and execute the command `no system-user linuxadmin disable`.

```
OS10(config)# no system-user linuxadmin disable
OS10(config)#
```

RADIUS authentication

To configure a RADIUS server for authentication, enter the server IP address or host name, and the key used to authenticate the OS10 switch on a RADIUS host. You can enter the authentication key in plain text or encrypted format. You can change the User Datagram Protocol (UDP) port number on the server.

- Configure a RADIUS authentication server in CONFIGURATION mode. By default, a RADIUS server uses UDP port 1812.

  ```
  radius-server host {hostname | ip-address}  key {0 authentication-key | 9 authentication-key
  | authentication-key} [auth-port port-number]
  ```

Re-enter the `radius-server host` command multiple times to configure more than one RADIUS server. If you configure multiple RADIUS servers, OS10 attempts to connect in the order you configured them. An OS10 switch connects with the configured RADIUS servers one at a time, until a RADIUS server responds with an accept or reject response. The switch tries to connect with a server for the configured number of retransmit retries and timeout period.

Configure global settings for the timeout and retransmit attempts allowed on RADIUS servers. By default, OS10 supports three RADIUS authentication attempts and times out after five seconds. No source interface is configured. The default VRF instance is used to contact RADIUS servers.

NOTE: You cannot configure both a non-default VRF instance and a source interface at the same time for RADIUS authentication.

- Configure the number of times OS10 retransmits a RADIUS authentication request in CONFIGURATION mode, from 0 to 100 retries; the default is 3.

  ```
  radius-server retransmit retries
  ```

- Configure the timeout period used to wait for an authentication response from a RADIUS server in CONFIGURATION mode, from 0 to 1000 seconds; the default is 5.

  ```
  radius-server timeout seconds
  ```

- (Optional) Specify an interface whose IP address is used as the source IP address for user authentication with RADIUS servers in CONFIGURATION mode. By default, no source interface is configured. OS10 selects the source IP address of any interface from which a packet is sent to a RADIUS server.

  NOTE: You must configure the IPv4 and all IPv6 addresses of the OS10 source interface in the client configuration on each RADIUS server. RADIUS authentication fails if the source-interface IP addresses are not configured on a RADIUS server.

  ```
  ip radius source-interface interface
  ```

- (Optional) By default, the switch uses the default VRF instance to communicate with RADIUS servers. You can optionally configure a non-default or the management VRF instance for RADIUS authentication in CONFIGURATION mode.

  ```
  radius-server vrf management
  radius-server vrf vrf-name
  ```

Configure RADIUS server

```
OS10(config)# radius-server host 1.2.4.5
OS10(config)# radius-server retransmit 10
OS10(config)# radius-server timeout 10
OS10(config)# ip radius source-interface mgmt 1/1/1
```

Configure RADIUS server for non-default VRFs

```
OS10(config)# ip vrf blue
OS10(conf-vrf)# exit
OS10(config)# radius-server vrf blue
```

View RADIUS server configuration

```
OS10# show running-configuration
...
radius-server host 1.2.4.5 key 9
3a95c26b2a5b96a6b80036839f296babe03560f4b0b7220d6454b3e71bdfc59b
radius-server retransmit 10
radius-server timeout 10
ip radius source-interface mgmt 1/1/1
...
```

Delete RADIUS server

```
OS10# no radius-server host 1.2.4.5
```

RADIUS over TLS authentication

Traditional RADIUS-based user authentication runs over UDP and uses the MD5 message-digest algorithm for secure communications. To provide enhanced security in RADIUS user authentication exchanges, RFC 6614 defines the RADIUS over Transport Layer Security (TLS) protocol. RADIUS over TLS secures the entire authentication exchange in a TLS connection and provides additional security by:

- Performing mutual authentication of a client and server using public key infrastructure (PKI) certificates
- Encrypting the entire authentication exchange so that neither user ID nor password is vulnerable to discovery

RADIUS over TLS authentication requires that X.509v3 PKI certificates are configured on a certification authority (CA) and installed on the switch. For more information, including a complete RADIUS over TLS use case, see X.509v3 certificates.

NOTE: RADIUS over TLS operates in FIPS mode when you enable FIPS using the `crypto fips enable` command. In FIPS mode, RADIUS over TLS requires that a FIPS-compliant certificate is installed on the switch. In non-FIPS mode, RADIUS over TLS requires that a certificate is installed as a non-FIPS certificate. For information about how to install FIPS-compliant and non-FIPS certificates, see Request and install host certificates.

To configure RADIUS over TLS user authentication, use the `radius-server host tls` command. Enter the server IP address or host name, and the shared secret key used to authenticate the OS10 switch on a RADIUS host. You must enter the name of an X.509v3 security profile to use with RADIUS over TLS authentication — see Security profiles. You can enter the authentication key in plain text or encrypted format. By default, RADIUS over TLS connections use TCP port 2083, and require that the authentication key is `radsec`. You can change the TCP port number on the server.

- Configure a RADIUS over TLS authentication on a RADIUS server in CONFIGURATION mode.

```
radius-server host {hostname | ip-address} tls security-profile profile-name
[auth-port port-number] key {0 authentication-key | 9 authentication-key | authentication-key}
```

To configure more than one RADIUS server for RADIUS over TLS authentication, re-enter the `radius-server host tls` command multiple times. If you configure multiple RADIUS servers, OS10 attempts to connect in the order you configured them. An OS10 switch connects with the configured RADIUS servers one at a time, until a RADIUS server responds with an accept or reject response. The switch tries to connect with a server for the configured number of retransmit retries and timeout period.

Enter the name of a security profile to use with RADIUS over TLS authentication. The security profile determines the X.509v3 certificate on the switch to use for TLS authentication with a RADIUS server. To configure a security profile for an OS10 application, see Security profiles.

Configure global settings for the timeout and retransmit attempts allowed on RADIUS servers as described in RADIUS authentication.

Configure RADIUS over TLS authentication server

```
OS10(config)# radius-server host 1.2.4.5 tls security-profile radius-prof key radsec
OS10(config)# radius-server retransmit 10
OS10(config)# radius-server timeout 10
```

TACACS+ authentication

Configure a TACACS+ authentication server by entering the server IP address or host name. You must also enter a text string for the key used to authenticate the OS10 switch on a TACACS+ host. The Transmission Control Protocol (TCP) port entry is optional.

TACACS+ provides greater data security by encrypting the entire protocol portion in a packet sent from the switch to an authentication server. RADIUS encrypts only passwords.

- Configure a TACACS+ authentication server in CONFIGURATION mode. By default, a TACACS+ server uses TCP port 49 for authentication.

  ```
  tacacs-server host {hostname | ip-address}  key {0 authentication-key | 9 authentication-key
  | authentication-key} [auth-port port-number]
  ```

  Re-enter the `tacacs-server host` command multiple times to configure more than one TACACS+ server. If you configure multiple TACACS+ servers, OS10 attempts to connect in the order you configured them. An OS10 switch connects with the configured TACACS + servers one at a time, until a TACACS+ server responds with an accept or reject response.

Configure global timeout setting allowed on TACACS+ servers. By default, OS10 times out after five seconds. No source interface is configured. The default VRF instance is used to contact TACACS+ servers.

NOTE: You cannot configure both a non-default VRF instance and a source interface at the same time for TACACS+ authentication.

- Configure the global timeout used to wait for an authentication response from TACACS+ servers in CONFIGURATION mode, from 1 to 1000 seconds; the default is 5.

  ```
  tacacs-server timeout seconds
  ```

- (Optional) Specify an interface whose IP address is used as the source IP address for user authentication with a TACACS+ server in CONFIGURATION mode. By default, no source interface is configured. OS10 selects the source IP address of any interface from which a packet is sent to a TACACS+ server.

  NOTE: If a TACACS+ server uses an IPv4 address, you must configure the IPv4 address of the OS10 source interface in the configuration file on the TACACS+ server. If a TACACS+ server uses an IPv6 address, you must configure all IPv6 addresses of the OS10 source interface in the TACACS+ configuration file. TACACS+ authentication fails if the source-interface IP address is not configured on a TACACS+ server.

  ```
  ip tacacs source-interface interface
  ```

- (Optional) By default, the switch uses the default VRF instance to communicate with TACACS+ servers. You can optionally configure a non-default or the management VRF instance for TACACS+ authentication in CONFIGURATION mode.

  ```
  tacacs-server vrf management
  tacacs-server vrf vrf-name
  ```

Configure TACACS+ server

```
OS10(config)# tacacs-server host 1.2.4.5 key mysecret
OS10(config)# ip tacacs source-interface loopback 2
```

Configure TACACS+ server for non-default VRFs

```
OS10(config)# ip vrf blue
OS10(conf-vrf)# exit
OS10(config)# tacacs-server vrf blue
```

View TACACS+ server configuration

```
OS10# show running-configuration
...
tacacs-server host 1.2.4.5 key 9
3a95c26b2a5b96a6b80036839f296babe03560f4b0b7220d6454b3e71bdfc59b
ip tacacs source-interface loopback 2
...
```

Delete TACACS+ server

```
OS10# no tacacs-server host 1.2.4.5
```

Unknown user role

When a RADIUS or TACACS+ server authenticates a user, it may return an unknown user role, or the role may be missing. In these cases, OS10 assigns the `netoperator` role and associated permissions to the user by default. You can reconfigure the default assigned role. In

addition, you can configure an unknown RADIUS or TACACS+ user-role name to inherit the permissions of an existing OS10 system-defined role.

- Reconfigure the default OS10 user role in CONFIGURATION mode.

  ```
  userrole {default | name} inherit existing-role-name
  ```

  - `default inherit` — Reconfigure the default permissions assigned to an authenticated user with a missing or unknown role.
  - *`name`* `inherit` — Enter the name of the RADIUS or TACACS+ user role that inherits permissions from an OS10 user role; 32 characters maximum.
  - *`existing-role-name`* — Assign the permissions associated with an existing OS10 user role:
    - `sysadmin` — Full access to all commands in the system, exclusive access to commands that manipulate the file system, and access to the system shell. A system administrator can create user IDs and user roles.
    - `secadmin` — Full access to configuration commands that set security policy and system access, such as password strength, AAA authorization, and cryptographic keys. A security administrator can display security information, such as cryptographic keys, login statistics, and log information.
    - `netadmin` — Full access to configuration commands that manage traffic flowing through the switch, such as routes, interfaces, and ACLs. A network administrator cannot access configuration commands for security features or view security information.
    - `netoperator` — Access only to EXEC mode to view the current configuration. A network operator cannot modify any configuration setting on a switch.

Reconfigure permissions for an unknown user role

```
OS10(config)# userrole default inherit sysadmin
```

Configure permissions for a RADIUS or TACACS+ user role

```
OS10(config)# userrole tacacsadmin inherit netadmin
```

SSH server

In OS10, the secure shell (SSH) server allows an SSH client to access an OS10 switch through a secure, encrypted connection. The SSH server authenticates remote clients using RADIUS challenge/response, a trusted host file, locally-stored passwords, and public keys.

Configure SSH server

- The SSH server is enabled by default. You can disable the SSH server using the `no ip ssh server enable` command.
- Challenge response authentication is disabled by default. To enable, use the `ip ssh server challenge-response-authentication` command.
- Host-based authentication is disabled by default. To enable, use the `ip ssh server hostbased-authentication` command.
- Password authentication is enabled by default. To disable, use the `no ip ssh server password-authentication` command.
- Public key authentication is enabled by default. To disable, use the `no ip ssh server pubkey-authentication` command.
- Password-less login is disabled by default. To enable, use the `username sshkey` or `username sshkey filename` commands.
- Configure the list of cipher algorithms using the `ip ssh server cipher` *`cipher-list`* command.
- Configure Key Exchange algorithms using the `ip ssh server kex` *`key-exchange-algorithm`* command.
- Configure hash message authentication code (HMAC) algorithms using the `ip ssh server mac` *`hmac-algorithm`* command.
- Configure the SSH server listening port using the `ip ssh server port` *`port-number`* command.
- Configure the SSH server to be reachable on the management VRF using the `ip ssh server vrf` command.
- Configure the SSH login timeout using the `ip ssh server login-grace-time` *`seconds`* command, from 0 to 300; default 60. To reset the default SSH prompt timer, use the `no ip ssh server login-grace-time` command.
- Configure the maximum number of authentication attempts using the `ip ssh server max-auth-tries` *`number`* command, from 0 to 10; default 6. To reset the default, use the `no ip ssh server max-auth-tries` command.

  The `max-auth-tries` value includes all authentication attempts, including public-key and password. If you enable both, public-key based authentication and password authentication, the public-key authentication is the default and is tried first. If it fails, the number of

`max-auth-tries` is reduced by one. In this case, if you configured `ip ssh server max-auth-tries 1`, the password prompt does not display.

Regenerate public keys

When enabled, the SSH server generates public keys by default and uses them for client authentication:

- A Rivest, Shamir, and Adelman (RSA) key using 2048 bits.
- An Elliptic Curve Digital Signature Algorithm (ECDSA) key using 256 bits
- An Ed25519 key using 256 bits

NOTE: RSA1 and DSA keys are not supported on the OS10 SSH server.

An SSH client must exchange the same public key to establish a secure SSH connection to the OS10 switch. If necessary, you can regenerate the keys used by the SSH server with a customized bit size. You cannot change the default size of the Ed25519 key. The `crypto key generate` command is available only to the `sysadmin` and `secadmin` roles.

1 Regenerate keys for the SSH server in EXEC mode.

```
crypto ssh-key generate {rsa {2048|3072|4096} | ecdsa {256|384|521} | ed25519}
```

2 Enter `yes` at the prompt to overwrite an existing key.

```
Host key already exists. Overwrite [confirm yes/no]:yes
Generated 2048-bit RSA key
```

3 Display the SSH public keys in EXEC mode.

```
show crypto ssh-key
```

After you regenerate SSH server keys, disable and re-enable the SSH server to use the new keys. Restarting the SSH server does not impact current OS10 sessions.

Virtual terminal line

Use Virtual terminal line (VTY) to control Telnet or SSH connections to the switch.

Enter VTY mode using the `line vty` command in CONFIGURATION mode.

```
OS10(config)# line vty
OS10(config-line-vty)#
```

Control access to VTY

You can control the Telnet or SSH connections to the switch by applying access lists on VTY lines.

Create IPv4 or IPv6 access lists with `permit` or `deny` filters.

Enter VTY mode using the `line vty` command in CONFIGURATION mode.

Apply the access lists to the VTY line with the `{ip | ipv6} access-class` *`access-list-name`* command.

Example

```
OS10(config)# ip access-list permit10
OS10(config-ipv4-acl)# permit ip 172.16.0.0 255.255.0.0 any
OS10(config-ipv4-acl)# exit
OS10(config)# line vty
OS10(config-line-vty)# ip access-class permit10
OS10(config-line-vty)#
```

View VTY ACL configuration

```
OS10(config-line-vty)# show configuration
!
```

```
line vty
 ip access-class permit10
 ipv6 access-class deny10
OS10(config-line-vty)#
```

Enable AAA accounting

To record information about all user-entered commands, use the AAA accounting feature — not supported for RADIUS accounting. AAA accounting records login and command information in OS10 sessions on console connections using the `console` option and remote connections using the `default` option, such as Telnet and SSH.

AAA accounting sends accounting messages:

- Sends a start notice when a process begins, and a stop notice when the process ends using the `start-stop` option
- Sends only a stop notice when a process ends using the `stop-only` option
- No accounting notices are sent using the `none` option
- Logs all accounting notices in syslog using the `logging` option
- Logs all accounting notices on configured TACACS+ servers using the `group tacacs+` option

Enable AAA accounting

- Enable AAA accounting in CONFIGURATION mode.

```
aaa accounting commands all {console | default} {start-stop | stop-only | none} [logging]
[group tacacs+]
```

The `no` version of this command disables AAA accounting.

Enable user lockout

By default, a maximum of three consecutive failed password attempts is supported on the switch. You can set a limit to the maximum number of allowed password retries with a specified lockout period for the user ID.

This feature is available only for the `sysadmin` and `secadmin` roles.

- Configure user lockout settings in CONFIGURATION mode.

```
password-attributes {[max-retry number ] [lockout-period minutes]}
```

  - `max-retry number` — Sets the maximum number of consecutive failed login attempts for a user before the user is locked out, from 0 to 16; default 3.
  - `lockout-period minutes` — Sets the amount of time that a user ID is prevented from accessing the system after exceeding the maximum number of failed login attempts, from 0 to 43,200; default 0.

When a user is locked out due to exceeding the maximum number of failed login attempts, other users can still access the switch.

By default, `lockout-period minutes` is 0; no lockout period is configured. Failed login attempts do not lock out a user.

Configure user lockout

```
OS10(config)# password-attributes max-retry 4 lockout period 360
```

Limit concurrent login sessions

To avoid an unlimited number of active sessions on a switch for the same user ID, you can limit the number of console and remote connections. Log in from a console connection by cabling a terminal emulator to the console serial port on the switch. Log in to the switch remotely through a virtual terminal line (VTY), such as Telnet and SSH.

- Configure the maximum number of concurrent login sessions in CONFIGURATION mode.
  ```
  OS10(config)# login concurrent-session limit number
  ```
  - `limit number` — Sets the maximum number of concurrent login sessions allowed for a user ID, from 1 to 12; default 10.

When you configure the maximum number of allowed concurrent login sessions, take into account that:

- Each remote VTY connection counts as one login session.
- All login sessions from a terminal emulator on an attached console count as one session.

Configure concurrent login sessions

```
OS10(config)# login concurrent-session limit 4
```

If you log in to the switch after the maximum number of concurrent sessions are active, an error message displays. To log in to the system, close one of your existing sessions.

```
OS10(config)# login concurrent-session limit 4

Too many logins for 'admin'.
Last login: Wed Jan 31 20:37:34 2018 from 10.14.1.213
Connection to 10.11.178.26 closed.
Current sessions for user admin:
Line              Location
2  vty 0          10.14.1.97
3  vty 1          10.14.1.97
4  vty 2          10.14.1.97
5  vty 3          10.14.1.97
```

Enable login statistics

To monitor system security, allow users to view their own login statistics when they sign in to the system. A large number of login failures or an unusual login location may indicate a system hacker. Enable the display of login information after a user successfully logs in; for example:

```
OS10 login: admin
Password:
Last login: Thu Nov  2 16:02:44 UTC 2017 on ttyS1
Linux OS10 3.16.43 #2 SMP Debian 3.16.43-2+deb8u5 x86_64
...
Time-frame for statistics     : 25 days
Role changed since last login : false
Failures since last login     : 0
Failures in time period       : 1
Successes in time period      : 14
OS10#
```

This feature is available only for the `sysadmin` and `secadmin` roles.

- Enable the display of login information in CONFIGURATION mode.
  ```
  login-statistics enable
  ```

To display information about user logins, use the `show login-statistics` command.

Enable login statistics

```
OS10(config)# login-statistics enable
```

To disable login statistics, use the `no login-statistics enable` command.

Privilege levels overview

Limiting access to the system is one method of protecting the system and your network. However, at times, you might need to provide others access to the router. To increase security, you can limit access to a subset of commands. OS10 enables you to configure a privilege level for users who need limited access to the system.

With the OS10, you can configure various privilege levels of various command modes. The system supports 16 privilege levels. The following lists the privilege levels:

- Level 0—Has the least privilege, provides access to basic commands.
- Level 1—The default configuration that provides access to a set of show commands and certain operations such as ping, traceroute, and so on.
- Level 15—Provides access to all available commands, equivalent to the `sysadmin` role.
- Levels 0, 1, and 15—Preconfigured by the system and are not user-configurable.
- Levels 2 to 14—Not configured You can customize these levels for different users and access rights.

Privilege levels inherit all permitted commands from all lower levels. For example, if you log in to privilege level 10 using the `enable 10` command, all commands that are assigned to privilege level 10 and lower are available for use.

You cannot configure a privilege level less than 2 for users assigned to the `sysadmin`, `netadmin`, and `secadmin` roles. You can configure users assigned to the `netoperator` role with privilege levels 0 or 1.

After you configure privilege levels, you can assign them to users using the `username` command or to EXEC mode using the `enable` command.

After you configure passwords for the EXEC mode levels, users can use the `enable privilege-level` command to switch between privilege levels. Using the `disable` command takes the user to a lower level.

When a user is assigned to a role, and is also assigned a privilege level, the user can access the commands available for that role and the commands available for that privilege level.

When a remote user logs in, OS10 checks for a match in the local system. If there is a local user as the remote user, the privilege level of the local user is applied to the remote user for the login session. If there is no match in the local system, depending on the role of the remote user, OS10 assigns default privilege levels. For `sysadmin`, `secadmin`, and `netadmin` roles, OS10 assigns level 15 and for the `netoperator` role, OS10 assigns level 1.

NOTE: The role of a local user and the corresponding remote user should be the same at both remote and local ends.

Configure privilege levels for users

To restrict CLI access for different users, create the required privilege levels, assign commands, and then assign privilege levels to users.

1 Configure privilege levels.

CONFIGURATION

```
privilege mode priv-lvl privilege-level command-string
```

- *mode*—Enter the privilege mode where you are configuring the specific command. The following table lists the available privilege modes and their corresponding command modes:

Privilege mode	CLI mode
Exec	exec
configure	class-map, DHCP, logging, monitor, openflow, policy-map, QOS, support-assist, telemetry, CoS, Tmap, UFD, VLT, VN, VRF, WRED, or alias
interface	Ethernet, FC, Loopback, mgmt, null, port-group, lag, breakout, range, port-channel, VLAN
route-map	route-map
router	router-bgp or router-ospf
line	line-vty

- `priv-lvl`—Enter the keyword and then the privilege number, from 2 to 14.
- *`command-string`*—Enter the specific command.

You cannot configure a privilege level less than 2 for users assigned to `sysadmin`, `netadmin`, and `secadmin` roles.

2 Create a user name and password and assign a privilege level.

CONFIGURATION

```
username username password password role role [ priv-lvl privilege-level]
```

- `username` *`username`*—Enter a text string. A maximum of 32 alphanumeric characters; one character minimum.
- `password` *`password`*—Enter a text string. A maximum of 32 alphanumeric characters; nine characters minimum.
- `role` *`role`*—Enter a user role:
  - `sysadmin`—Full access to all commands in the system, exclusive access to commands that manipulate the file system, and access to the system shell. A system administrator can create user IDs and user roles.
  - `secadmin`—Full access to configuration commands that set security policy and system access, such as password strength, AAA authorization, and cryptographic keys. A security administrator can display security information, such as cryptographic keys, login statistics, and log information.
  - `netadmin`—Full access to configuration commands that manage traffic flowing through the switch, such as routes, interfaces, and access control lists (ACLs). A network administrator cannot access configuration commands for security features or view security information.
  - `netoperator`—Access to EXEC mode to view the current configuration. A network operator cannot modify any configuration setting on a switch.
- `priv-lvl` *`privilege-level`*—Enter a privilege level, from 0 to 15. If you do not specify the `priv-lvl` option, the system assigns privilege level 1 for the `netoperator` user and privilege level 15 for the `sysadmin`, `secadmin`, and `netadmin` users.

The following is an example of configuring privilege levels and assigning them to a user:

```
OS10(config)# privilege exec priv-lvl 12 "show version"
OS10(config)# privilege exec priv-lvl 12 "configure terminal"
OS10(config)# privilege configure priv-lvl 12 "interface ethernet"
OS10(config)# privilege interface priv-lvl 12 "ip address"
OS10(config)# username delluser password $6$Yij02Phe2n6whp7b$ladskj0HowijIlkajg981 role
secadmin priv-lvl 12
```

The following example shows the privilege level of the current user:

```
OS10# show privilege
Current privilege level is 15.
```

The following example displays the privilege levels of all users who are logged into OS10:

```
OS10# show users

Index Line   User  Role     Application Idle Login-Time            Location         Privilege
----- ------ ----- -----   ----------- ---- -----------            --------         ---------
1     pts/0  admin sysadmin bash        >24h 2018-09-08 T06:51:37Z 10.14.1.91 [ssh] 15
2     pts/1  netad netadmin bash        >24h 2018-09-08 T06:54:33Z 10.14.1.91 [ssh] 10
```

Configure enable password

To configure an enable password:

1 Configure a privilege level and assign commands to it.

CONFIGURATION

```
privilege mode priv-lvl privilege-level command-string
```

- *mode*—Enter the privilege mode where you are configuring the specific command. The following table lists the available privilege modes and their corresponding command modes:

Privilege mode	CLI mode
Exec	exec
configure	class-map, DHCP, logging, monitor, openflow, policy-map, QOS, support-assist, telemetry, CoS, Tmap, UFD, VLT, VN, VRF, WRED, or alias
interface	Ethernet, FC, loopback, mgmt, null, port-group, lag, breakout, range, port-channel, VLAN
route-map	route-map
router	router-bgp, router-ospf
line	line-vty

- `priv-lvl`—Enter the keyword and then the privilege number, from 2 to 14.
- *command-string*—Enter the specific command.

You cannot configure a privilege level less than 2 for `sysadmin`, `netadmin`, and `secadmin` roles.

2 Configure an enable password and assign the privilege level to it.

CONFIGURATION

```
enable password encryption-type password-string priv-lvl privilege-level
```

Encryption types:

- `0`—Specifies an unencrypted password follows
- `sha-256`—Specifies a SHA-256 encrypted password follows
- `sha-512`—Specifies a SHA-512 encrypted password follows

`priv-lvl`—Enter the keyword and then the privilege number. The range is from 1to 15.

```
OS10(config)# privilege exec priv-lvl 3 "show version"
OS10(config)# enable password 0 P@$$w0Rd priv-lvl 3
```

```
OS10(config)# privilege exec priv-lvl 12 "configure terminal"
OS10(config)# privilege configure priv-lvl 12 route-map
OS10(config)# privilege route-map priv-lvl 12 "set local-preference"
OS10(config)# enable password sha-256 $5$2uThib1o$84p.tykjmz/w7j26ymoKBjrb7uepkUB priv-lvl 12
```

Audit log

To monitor user activity and configuration changes on the switch, enable the audit log. Only the `sysadmin` and `secadmin` roles can enable, view, and clear the audit log.

The audit log records configuration and security events, including:

- User logins and logouts on the switch, failed logins, and concurrent login attempts by a user

- User-based configuration changes recorded with the user ID, date, and time of the change. The specific configuration parameters that were changed are not logged.
- Establishment of secure traffic flows, such as SSH, and violations on secure flows
- Certificate issues, including user access and changes made to certificate installation using `crypto` commands
- Adding and deleting users

Audit log entries are saved locally and sent to configured Syslog servers. To set up a Syslog server, see System logging.

Enable audit log

- Enable the recording of configuration and security events in the audit log on Syslog servers in CONFIGURATION mode.

```
logging audit enable
```

To disable audit logging, enter the `no logging audit enable` command.

View audit log

- Display audit log entries in EXEC mode. By default, 24 entries are displayed, starting with the oldest event. Enter `reverse` to display entries starting with the most recent events. You can change the number of entries displayed.

```
show logging audit [reverse] [number]
```

Clear audit log

- Clear all events in the audit log in CONFIGURATION mode.

```
clear logging audit
```

Example

```
OS10(config)# logging audit enable
OS10(config)# exit

OS10# show logging audit 4
<14>1 2019-02-14T13:15:06.283337+00:00 OS10 audispd - - - Node.1-Unit.1:PRI [audit], Dell EMC
(OS10)  node=OS10 type=USER_END msg=audit(1550150106.277:597): pid=7908 uid=0 auid=4294967295
ses=4294967295 msg='op=PAM:session_close acct="admin" exe="/bin/su" hostname=? addr=?
terminal=??? res=success'
<110>1 2019-02-14T13:15:16.331515+00:00 OS10 .clish 7412 - -  Node.1-Unit.1:PRI [audit], User
admin on console used cmd: 'crypto security-profile mltestprofile' - success
<110>1 2019-02-14T13:15:21.794529+00:00 OS10 .clish 7412 - -  Node.1-Unit.1:PRI [audit], User
admin on console used cmd: 'exit' - success
<110>1 2019-02-14T13:16:05.882555+00:00 OS10 .clish 7412 - -  Node.1-Unit.1:PRI [audit], User
admin on console used cmd: 'exit' - success

OS10# show logging audit reverse 4
<110>1 2019-02-14T13:16:05.882555+00:00 OS10 .clish 7412 - -  Node.1-Unit.1:PRI [audit], User
admin on console used cmd: 'exit' - success
<110>1 2019-02-14T13:15:21.794529+00:00 OS10 .clish 7412 - -  Node.1-Unit.1:PRI [audit], User
admin on console used cmd: 'exit' - success
<110>1 2019-02-14T13:15:16.331515+00:00 OS10 .clish 7412 - -  Node.1-Unit.1:PRI [audit], User
admin on console used cmd: 'crypto security-profile mltestprofile' - success
<14>1 2019-02-14T13:15:06.283337+00:00 OS10 audispd - - - Node.1-Unit.1:PRI [audit], Dell EMC
(OS10)  node=OS10 type=USER_END msg=audit(1550150106.277:597): pid=7908 uid=0 auid=4294967295
ses=4294967295 msg='op=PAM:session_close acct="admin" exe="/bin/su" hostname=? addr=?
terminal=??? res=success'OS10# show logging audit reverse 10
```

Security commands

aaa accounting

Enables AAA accounting.

Syntax

```
aaa accounting commands all {console | default} {start-stop | stop-only | none}
[logging] [group tacacs+]
```

Parameters

- `commands all` — Record all user-entered commands. This option is not supported for RADIUS accounting.
- `console` — Record all user authentication and logins or all user-entered commands in OS10 sessions on console connections.
- `default` — Record all user authentication and logins or all user-entered commands in OS10 sessions on remote connections; for example, Telnet and SSH.
- `start-stop` — Send a start notice when a process begins, and a stop notice when the process ends.
- `stop-only` — Send only a stop notice when a process ends.
- `none` — No accounting notices are sent.
- `logging` — Logs all accounting notices in syslog.
- `group tacacs+` — Logs all accounting notices on the first reachable TACACS+ server.

Default AAA accounting is disabled.

Command Mode CONFIGURATION

Usage Information You can enable the recording of accounting events in both the syslog and on TACACS+ servers. The `no` version of the command disables AAA accounting.

Example

```
OS10(config)# aaa accounting commands all console start-stop logging group
tacacs+
```

Supported Releases 10.4.1.0 or later

aaa authentication login

Configures the AAA authentication method used for console, and SSH and Telnet logins.

Syntax

```
aaa authentication login {console | default} {local | group radius | group
tacacs+}
```

Parameters

- `console` — Configure authentication methods for console logins.
- `default` — Configure authentication methods for SSH and Telnet logins.
- `local` — Use the local username, password, and role entries configured with the `username password role` command.
- `group radius` — Use the RADIUS servers configured with the `radius-server host` command.
- `group tacacs+` — Use the TACACS+ servers configured with the `tacacs-server host` command.

Default Local authentication

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes all configured authentication methods and defaults to using local authentication.

Example

```
OS10(config)# aaa authentication login default group radius local
OS10(config)# do show running-configuration aaa
```

```
aaa authentication login default group radius local
aaa authentication login console local
```

```
OS10(config)# no aaa authentication login default
OS10(config)# do show running-configuration aaa
aaa authentication login default local
aaa authentication login console local
```

Supported Releases 10.4.1.0 or later

aaa re-authenticate enable

Requires user re-authentication after a change in the authentication method or server.

Syntax `aaa re-authenticate enable`

Parameters None

Default Disabled

Command Mode EXEC

Usage Information

- After you enable user re-authentication and change the authentication method or server, users are logged out of the switch and are prompted to log in again to re-authenticate. User re-authentication is triggered by:
  - Adding or removing a RADIUS server as a configured server host with the `radius-server host` command.
  - Adding or removing an authentication method with the `aaa authentication [local | radius]` command.
- The `no` version of the command disables user re-authentication.

Example

```
OS10(config)# aaa re-authenticate enable
```

Supported Releases 10.4.0E(R1) or later

boot protect disable username

Allows you to disable bootloader protection.

Syntax `boot protect disable username` *username*

Parameters None

Default Disabled

Command Mode EXEC

Usage Information You can disable bootloader protection for each individual user.

Example

```
OS10# boot protect disable username root
```

Supported Releases 10.4.3.0 or later

boot protect enable username password

Allows you to enable bootloader protection.

Syntax `boot protect enable username username password password`

Parameters

- *username* — Enter the username to provide access to bootloader protection.
- *password* — Enter a password for the specified username.

Default Disabled

Command Mode EXEC

Usage Information You can enable bootloader protection by executing this command. You can configure a maximum of three username / password pairs for bootloader protection.

Example

```
OS10# boot protect enable username root password calvin
```

Supported Releases 10.4.3.0 or later

clear logging audit

Deletes all events in the audit log.

Syntax `clear logging audit`

Parameters None

Defaults Not configured

Command Mode EXEC

Usage Information To display the contents of the audit log, use the `show logging audit` command.

Example

```
OS10# clear logging audit
Proceed to clear all audit log messages [confirm yes/no(default)]:yes
```

Supported Releases 10.4.3.0 or later

crypto ssh-key generate

Regenerate public keys used in SSH authentication.

Syntax `crypto ssh-key generate {rsa bits | ecdsa bits | ed25519}`

Parameters

- `rsa bits` — Regenerates the RSA key with the specified bit size (2048, 3072, or 4096; default 2048).
- `ecdsa bits` — Regenerates the ECDSA key with the specified bit size (256, 384, or 521; default 256).
- `ed25519` — Regenerates the Ed25519 key with the default bit size.

Default The SSH server uses default public key lengths for client authentication:

- RSA key: 2048 bits

- ECDSA key : 256 bits
- Ed25519 key: 256 bits

Command Mode EXEC

Usage Information If necessary, you can regenerate the public keys used by the SSH server with a customized bit size. You cannot change the default size of the Ed25519 key. The `crypto ssh-key generate` command is available only to the `sysadmin` and `secadmin` roles.

Example

```
OS10# crypto ssh-key generate rsa 4096
Host key already exists. Overwrite [confirm yes/no]:yes
Generated 4096-bit RSA key
OS10#
```

Supported Releases 10.4.1.0 or later

disable

Exits Privilege Exec mode or lowers privilege level.

Syntax `disable` *`privilege-level`*

Parameters

- *`privilege-level`*—Enter the configured privilege level, from 0 to 15.

Defaults 1

Command Mode Privileged EXEC

Usage Information If you do not specify a privilege level, the system assigns level 1.

Example

```
OS10# disable
```

```
OS10# disable 6
```

Supported Releases 10.4.3.0 or later

enable

Enables Exec Privilege mode or a specific privilege level.

Syntax `enable` *`privilege-level`*

Parameters

- *`privilege-level`*—Enter the configured privilege level, from 0 to 15.

Defaults 15

Command Mode Exec

Usage Information Dell EMC Networking recommends configuring a password for privilege level 15 using the `enable password` command. If you do not configure a password for a level, you can switch to that level without entering a password, unless a password is configured for a highest intermediate level. If you configure a password for an intermediate level, you can enter that password. To access privilege level 15, you must configure the `enable password` command. If you do not configure a password for privilege level 15, you cannot enter level 15. For privilege levels 0 to 14, the `enable password` command is optional.

Privilege levels inherit all permitted commands from all lower levels. For example, if you log in to privilege level 10 using the `enable 10` command, all commands that are assigned to privilege level 10 and lower are available for use.

Example

```
OS10# enable
OS10# enable 10
```

Supported Releases 10.4.3.0 or later

enable password

Set a password for a specific privilege level.

Syntax `enable password encryption-type password-string priv-lvl privilege-level`

Parameters

- *encryption-type*—Enter the encryption type. The system supports the following encryption types:
  - `0`—Specifies an unencrypted password follows
  - `sha-256`—Specifies a SHA-256 encrypted password follows
  - `sha-512`—Specifies a SHA-512 encrypted password follows
- `priv-lvl`—Enter the keyword and then the privilege number, from 1 to 15.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes a privilege.

Example

```
OS10(conf)# enable password 0 P@$$w0Rd priv-lvl 12
OS10(conf)# enable password sha-256 $5$2uThib1o$84p.tykjmz/
w7j26ymoKBjrb7uepkUB priv-lvl 12
OS10(conf)# enable password sha-512 $6$Yij02Phe2n6whp7b$ladskj0HowijIlkajg981
priv-lvl 12
```

Supported Releases 10.4.3.0 or later

ip access-class

Filters connections based on an IPv4 access list in virtual terminal line.

Syntax `ip access-class access-list-name`

Parameters *access-list-name*—Enter the access list name.

Default Not configured

Command Mode LINE VTY CONFIGURATION

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# line vty
OS10(config-line-vty)# ip access-class deny10
```

Supported Releases 10.4.0E(R1) or later

ip radius source-interface

Specifies the interface whose IP address is used as the source IP address for user authentication with a RADIUS server.

Syntax `ip radius source-interface interface`

Parameters *interface*:

- `ethernet node/slot/port[:subport]` — Enter a physical Ethernet interface.
- `loopback number` — Enter a Loopback interface, from 0 to 16383.
- `mgmt 1/1/1` — Enter the management interface.
- `port-channel channel-id` — Enter a port-channel ID, from 1 to 28.
- `vlan vlan-id` — Enter a VLAN ID, from 1 to 4093.

Default Not configured.

Command Mode CONFIGURATION

Usage Information By default, no source interface is configured. OS10 selects the source IP address as the IP address of the interface from which a packet is sent to the RADIUS server. The `no` version of this command removes the configured source interface.

Example

```
OS10(config)# ip radius source-interface ethernet 1/1/10
```

Supported Releases 10.4.3.0 or later

ip tacacs source-interface

Specifies the interface whose IP address is used as the source IP address for user authentication with a TACACS+ server.

Syntax `ip tacacs source-interface interface`

Parameters *interface*:

- `ethernet node/slot/port[:subport]` — Enter a physical Ethernet interface.
- `loopback number` — Enter a Loopback interface, from 0 to 16383.
- `mgmt 1/1/1` — Enter the management interface.
- `port-channel channel-id` — Enter a port-channel ID, from 1 to 28.
- `vlan vlan-id` — Enter a VLAN ID, from 1 to 4093.

Default Not configured.

Command Mode CONFIGURATION

Usage Information By default, no source interface is configured. OS10 selects the source IP address as the IP address of the interface from which a packet is sent to the TACACS+ server. The `no` version of this command removes the configured source interface.

Example

```
OS10(config)# ip tacacs source-interface ethernet 1/1/10
```

Supported Releases 10.4.1.0 or later

ipv6 access-class

Filters connections based on an IPv6 access list in virtual terminal line.

Syntax	`ipv6 access-class` *`access-list-name`*
Parameters	*`access-list-name`*—Enter the access list name.
Default	Not configured
Command Mode	LINE VTY CONFIGURATION
Usage Information	The `no` version of this command removes the filter.
Example	`OS10(config)# line vty` `OS10(config-line-vty)# ipv6 access-class permit10`
Supported Releases	10.4.0E(R1) or later

ip ssh server challenge-response-authentication

Enable challenge response authentication in an SSH server.

Syntax	`ip ssh server challenge-response-authentication`
Parameters	None
Default	Disabled
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command disables the challenge response authentication.
Example	`OS10(config)# ip ssh server challenge-response-authentication`
Supported Releases	10.3.0E or later

ip ssh server cipher

Configure the list of cipher algorithms in the SSH server.

Syntax	`ip ssh server cipher` *`cipher-list`*
Parameters	*`cipher-list`* — Enter the list of cipher algorithms separated by space. The following is the list of cipher algorithms the SSH server supports: • 3des-cbc • aes128-cbc • aes192-cbc • aes256-cbc • aes128-ctr • aes192-ctr • aes256-ctr • aes128-gcm@openssh.com • aes256-gcm@openssh.com

- blowfish-cbc
- cast128-cbc
- chacha20-poly1305@opens

Default

- aes128-ctr
- aes192-ctr
- aes256-ctr
- aes128-gcm@openssh.com
- aes256-gcm@openssh.com
- chacha20-poly1305@opens

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the configuration.

Example

```
OS10(config)# ip ssh server cipher 3des-cbc aes128-cbc
```

Supported Releases 10.3.0E or later

ip ssh server enable

Enable the SSH server.

Syntax

```
ip ssh server enable
```

Parameters None

Default Enabled

Command Mode CONFIGURATION

Usage Information The `no` version of this command disables the SSH server.

Example

```
OS10(config)# ip ssh server enable
```

Supported Releases 10.3.0E or later

ip ssh server hostbased-authentication

Enable host-based authentication in an SSH server.

Syntax

```
ip ssh server hostbased-authentication
```

Parameters None

Default Disabled

Command Mode CONFIGURATION

Usage Information The `no` version of this command disables the host-based authentication.

Example

```
OS10(config)# ip ssh server hostbased-authentication
```

Supported Releases 10.3.0E or later

ip ssh server kex

Configure the list of Key Exchange algorithms in the SSH server.

Syntax `ip ssh server kex key-exchange-algorithm`

Parameters *key-exchange-algorithm* — Enter the list of Key Exchange algorithms separated by space. The following is the list of Key Exchange algorithms the SSH server supports:

- curve25519-sha256
- diffie-hellman-group1-sha1
- diffie-hellman-group14-sha1
- diffie-hellman-group-exchange-sha1
- diffie-hellman-group-exchange-sha256
- ecdh-sha2-nistp256
- ecdh-sha2-nistp384
- ecdh-sha2-nistp521

Default

- curve25519-sha256
- diffie-hellman-group14-sha1
- diffie-hellman-group-exchange-sha256
- ecdh-sha2-nistp256
- ecdh-sha2-nistp384
- ecdh-sha2-nistp521

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the configuration.

Example

```
OS10(config)# ip ssh server kex curve25519-sha256 diffie-hellman-group1-sha1
```

Supported Releases 10.3.0E or later

ip ssh server mac

Configure the list of hash message authentication code (HMAC) algorithms in the SSH server.

Syntax `ip ssh server mac hmac-algorithm`

Parameters *hmac-algorithm* — Enter the list of HMAC algorithms separated by space. The following is the list of HMAC algorithms the SSH server supports:

- hmac-md5
- hmac-md5-96
- hmac-ripemd160
- hmac-sha1
- hmac-sha1-96
- hmac-sha2-256
- hmac-sha2-512
- umac-64@openssh.com

- umac-128@openssh.com
- hmac-md5-etm@openssh.com
- hmac-md5-96-etm@openssh.com
- hmac-ripemd160-etm@openssh.com
- hmac-sha1-etm@openssh.com
- hmac-sha1-96-etm@openssh.com
- hmac-sha2-256-etm@openssh.com
- hmac-sha2-512-etm@openssh.com
- umac-64-etm@openssh.com
- umac-128-etm@openssh.com

Default

- hmac-sha1
- hmac-sha2-256
- hmac-sha2-512
- umac-64@openssh.com
- umac-128@openssh.com
- hmac-sha1-etm@openssh.com
- hmac-sha2-256-etm@openssh.com
- hmac-sha2-512-etm@openssh.com
- umac-64-etm@openssh.com
- umac-128-etm@openssh.com

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the configuration.

Example

```
OS10(config)# ip ssh server mac hmac-md5 hmac-md5-96 hmac-ripemd160
```

Supported Releases 10.3.0E or later

ip ssh server password-authentication

Enable password authentication in an SSH server.

Syntax

```
ip ssh server password-authentication
```

Parameters None

Default Enabled

Command Mode CONFIGURATION

Usage Information The `no` version of this command disables the password authentication.

Example

```
OS10(config)# ip ssh server password-authentication
```

Supported Releases 10.3.0E or later

ip ssh server port

Configure the SSH server listening port.

Syntax	`ip ssh server port port-number`
Parameters	*port-number* — Enter the listening port number, from 1 to 65535.
Default	22
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command removes the configuration.
Example	`OS10(config)# ip ssh server port 255`
Supported Releases	10.3.0E or later

ip ssh server pubkey-authentication

Enable public key authentication in an SSH server.

Syntax	`ip ssh server pubkey-authentication`
Parameters	None
Default	Enabled
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command disables the public key authentication.
Example	`OS10(config)# ip ssh server pubkey-authentication`
Supported Releases	10.3.0E or later

ip ssh server vrf

Configures an SSH server for the management or non-default VRF instance.

Syntax	`ip ssh server vrf {management | vrf vrf-name}`
Parameters	• `management` — Configures the management VRF instance to reach the SSH server. • `vrf vrf-name` — Enter the keyword `vrf` followed by the name of the VRF to configure that non-default VRF instance tp reach the SSH server.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	By default, the SSH server is enabled.
Example	`OS10(config)# ip ssh server vrf management` `OS10(config)# ip ssh server vrf vrf-blue`
Supported Releases	10.4.0E(R1) or later

line vty

Enters the virtual terminal line mode to access the virtual terminal (VTY).

Syntax	`line vty`
Parameters	None
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	None
Example	`OS10(config)# line vty` `OS10(config-line-vty)#`
Supported Releases	10.4.0E(R1) or later

logging audit enable

Enable the recording of configuration and security events in the audit log.

Syntax	`logging audit enable`
Parameters	None
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	Audit log entries are saved locally and sent to configured Syslog servers. Only the `sysadmin` and `secadmin` roles can enable the audit log. The `no` version of the command disables audit log recording.
Example	`OS10(conf)# logging audit enable`
Supported Releases	10.4.3.0 or later

login concurrent-session limit

Configures the maximum number of concurrent login sessions allowed for a user ID.

Syntax	`login concurrent-session limit` *number*
Parameters	limit *number* — Enter the limit of concurrent login sessions, from 1 to 12.
Default	10 concurrent login sessions are supported.
Command Mode	CONFIGURATION
Usage Information	The total number of concurrent login sessions for the same user ID includes all console and remote connections, where: • Each remote VTY connection counts as one login session. • All login sessions from a terminal emulator on an attached console count as one session. The `no` version of the command disables the configured number of allowed login sessions.

Example

```
OS10(config)# login concurrent-session limit 7
```

Supported Releases 10.4.1.0 or later

login-statistics enable

Enables the display of login statistics to users.

Syntax `login-statistics enable`

Parameters None

Default Disabled

Command Mode CONFIGURATION

Usage Information Only the `sysadmin` and `secadmin` roles have access to this command. When enabled, user login information, including the number of successful and failed logins, role changes, and the last time a user logged in, displays after a successful login. The `no login-statistics enable` command disables login statistics.

Example

```
OS10(config)# login-statistics enable
```

Supported Releases 10.4.0E(R1) or later

password-attributes

Configures rules for password entries.

Syntax `password-attributes {[min-length number] [character-restriction {[upper number] [lower number] [numeric number] [special-char number]}}`

Parameters

- `min-length number` — (Optional) Sets the minimum number of required alphanumeric characters, from 6 to 32; default 9.
- `character-restriction`:
  - `upper number` — (Optional) Sets the minimum number of uppercase characters required, from 0 to 31; default 0.
  - `lower number` — (Optional) Sets the minimum number of lowercase characters required, from 0 to 31; default 0.
  - `numeric number` — (Optional) Sets the minimum number of numeric characters required, from 0 to 31; default 0.
  - `special-char number` — (Optional) Sets the minimum number of special characters required, from 0 to 31; default 0.

Default

- Minimum length: 9 characters
- Uppercase characters: 0
- Lowercase characters: 0
- Numeric characters: 0
- Special characters: 0

Command Mode EXEC

Usage Information

- By default, the password you configure with the `username password` command must be at least nine alphanumeric characters.
- Use this command to increase password strength. When you enter the command, at least one parameter is required. When you enter the `character-restriction` parameter, at least one option is required.
- To reset parameters to their default values, enter the `no password-attributes` command.

Example

```
OS10(config)# password-attributes min-length 6 character-restriction upper 2
lower 2 numeric 2
```

Supported Releases 10.4.0E(R1) or later

password-attributes max-retry lockout-period

Configures a maximum number of consecutive failed login attempts and the lockout period for the user ID.

Syntax `password-attributes {[max-retry number] [lockout-period minutes]}`

Parameters

- `max-retry number` — (Optional) Sets the maximum number of consecutive failed login attempts for a user before the user is locked out, from 0 to 16.
- `lockout-period minutes` — (Optional) Sets the amount of time that a user ID is prevented from accessing the system after exceeding the maximum number of failed login attempts, from 0 to 43,200.

Default

- Maximum retries: 3 — A maximum of three failed login attempts is supported.
- Lockout period: 0 — No lockout period is configured. Failed login attempts do not lock out a user.

Command Mode CONFIGURATION

Usage Information

- To remove the configured `max-retry` or `lockout-period` settings, enter the `no password-attributes {max-retry | lockout-period}` command.
- When a user is locked out due to exceeding the maximum number of failed login attempts, other users can still access the switch.

Example

```
OS10(config)# password-attributes max-retry 5 lockout-period 30
```

Supported Releases 10.4.1.0 or later

privilege

Create a privilege level and associate commands with it.

Syntax `privilege mode priv-lvl privilege-level command-string`

Parameters

- *mode*—Enter the privilege mode where you are configuring the specific command. The following table lists the available privilege modes and their corresponding command modes:

Privilege mode	CLI mode
Exec	Exec
configure	class-map, DHCP, logging, monitor, openflow, policy-map, QOS, support-assist, telemetry, CoS, Tmap, UFD, VLT, VN, VRF, WRED, or alias

Privilege mode	CLI mode
interface	Ethernet, FC, loopback, mgmt, null, port-group, lag, breakout, range, port-channel, VLAN
route-map	route-map
router	router-bgp or router-ospf
line	line-vty

- `priv-lvl`—Enter the keyword and then the privilege number, from 2 to 14.
- *`command-string`*—Enter the specific command.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information You cannot configure a privilege level less than 2 for users assigned to `sysadmin`, `netadmin`, and `secadmin` roles.

If a command that you associate with a privilege level has a space, enter that command within double quotes. If a command does not have a space or separated by a hyphen, double quotes are not needed.

The `no` version of this command removes a command from a privilege level.

Example

```
OS10(config)# privilege exec priv-lvl 3 "configure terminal"
OS10(config)# privilege configure priv-lvl 3 "interface ethernet"
OS10(config)# privilege interface priv-lvl  "ip address"
OS10(config)# privilege configure priv-lvl 3 route-map
OS10(config)# privilege route-map priv-lvl 3 "set local-preference"
```

Supported Releases 10.4.3.0 or later

radius-server host

Configures a RADIUS server and the key used to authenticate the switch on the server.

Syntax

```
radius-server host {hostname | ip-address} key {0 authentication-key | 9
authentication-key | authentication-key} [auth-port port-number]
```

Parameters

- *`hostname`* — Enter the host name of the RADIUS server.
- *`ip-address`* — Enter the IPv4 (A.B.C.D) or IPv6 (x:x:x:x::x) address of the RADIUS server.
- `key 0` *`authentication-key`* — Enter an authentication key in plain text. A maximum of 42 characters.
- `key 9` *`authentication-key`* — Enter an authentication key in encrypted format. A maximum of 128 characters.
- *`authentication-key`* — Enter an authentication in plain text. A maximum of 42 characters. It is not necessary to enter `0` before the key.
- `auth-port` *`port-number`* — (Optional) Enter the UDP port number used on the server for authentication, from 0 to 65535, default 1812.

Default Not configured

Command Mode CONFIGURATION

Usage Information The authentication key must match the key configured on the RADIUS server. You cannot enter spaces in the key. The `show running-configuration` output displays both unencrypted and encrypted keys in encrypted format. Configure global settings for the timeout and retransmit attempts allowed on RADIUS servers using the

`radius-server retransmit` and `radius-server timeout` commands. The `no` version of this command removes a RADIUS server configuration.

Example

```
OS10(config)# radius-server host 1.5.6.4 key secret1
```

Supported Releases 10.2.0E or later

radius-server host tls

Configures a RADIUS server for RADIUS over TLS user authentication and secure communication. The `radsec` shared key and a security profile that uses an X.509v3 certificate is required for RADIUS over TLS authentication.

Syntax

```
radius-server host {hostname | ip-address} tls security-profile profile-name
[auth-port tcp-port-number] key {0 authentication-key | 9 authentication-key |
authentication-key}
```

Parameters

- *hostname* — Enter the host name of the RADIUS server.
- *ip-address* — Enter the IPv4 (A.B.C.D) or IPv6 (x:x:x:x::x) address of the RADIUS server.
- `tls` — Enter `tls` to secure RADIUS server communication using the TLS protocol.
- `security-profile` *profile-name* — Enter the name of an X.509v3 security profile to use with RADIUS over TLS authentication. To configure a security profile for an OS10 application, see Security profiles.
- `auth-port` *tcp-port-number* — (Optional) Enter the TCP port number that the server uses for authentication. The range is from 0 to 65535. The default is 2083.
- `key 0` *authentication-key* — Enter the `radsec` shared key in plain text.
- `key 9` *authentication-key* — Enter the `radsec` shared key in encrypted format.
- *authentication-key* — Enter the `radsec` shared key in plain text. It is not necessary to enter `0` before the key.

Default TCP port 2083 on a RADIUS server is used for RADIUS over TLS communication.

Command Mode CONFIGURATION

Usage Information For RADIUS over TLS authentication, configure the `radsec` shared key on the server and OS10 switch. The `show running-configuration` output displays both the unencrypted and encrypted key in encrypted format. Configure global settings for the timeout and retransmit attempts allowed on a RADIUS over TLS servers using the `radius-server retransmit` and `radius-server timeout` commands. The `no` version of this command removes a RADIUS server from RADIUS over TLS communication.

RADIUS over TLS authentication requires that X.509v3 PKI certificates are configured on a certification authority and installed on the switch. For more information, including a complete RADIUS over TLS example, see xref="X.509v3 certificates".

Example

```
OS10(config)# radius-server host 1.5.6.4 tls security-profile radius-admin key
radsec
```

Supported Releases 10.4.3.0 or later

radius-server retransmit

Configures the number of authentication attempts allowed on RADIUS servers.

Syntax

```
radius-server retransmit retries
```

Parameters *retries* — Enter the number of retry attempts, from 0 to 100.

Default	An OS10 switch retransmits a RADIUS authentication request three times.
Command Mode	CONFIGURATION
Usage Information	Use this command to globally configure the number of retransmit attempts allowed for authentication requests on RADIUS servers. The `no` version of this command resets the value to the default.
Example	`OS10(config)# radius-server retransmit 50`
Supported Releases	10.2.0E or later

radius-server timeout

Configures the timeout used to resend RADIUS authentication requests.

Syntax	`radius-server timeout` *seconds*
Parameters	*seconds* — Enter the time in seconds for retransmission, from 0 to 1000.
Default	An OS10 switch stops sending RADIUS authentication requests after five seconds.
Command Mode	CONFIGURATION
Usage Information	Use this command to globally configure the timeout value used on RADIUS servers. The `no` version of this command resets the value to the default.
Example	`OS10(config)# radius-server timeout 360`
Supported Releases	10.2.0E or later

radius-server vrf

Configures the RADIUS server for the management or non-default VRF instance.

Syntax	`radius-server vrf {management |` *vrf-name*`}`
Parameters	• `management` — Enter the keyword to configure the RADIUS server for the management VRF instance. • *vrf-name* — Enter the keyword then the name of the VRF to configure the RADIUS server for that non-default VRF instance.
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	Use this command to associate RADIUS servers with a VRF. If you do not configure a VRF on the RADIUS server list, the servers are on the default VRF. RADIUS server lists and VRFs have one-to-one mapping. When you remove the VRF, the RADIUS server lists are also removed automatically. The `no` version of this command removes the RADIUS server from the management VRF instance.
Example	`OS10(config)# radius-server vrf management` `OS10(config)# radius-server vrf blue`
Supported Releases	10.4.0E(R1) or later

show boot protect

Displays the current list of configured users that have access to bootloader protection.

Syntax `show boot protect`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information Displays the current list of authorised users for bootloader protection, but hides their passwords for security reasons.

Example (Disabled)

```
OS10# show boot protect
Boot protection disabled
```

Example (Enabled)

```
OS10# show boot protect
Boot protection enabled
Authorized users: root linuxadmin admin
```

Supported Releases 10.4.3.0 or later

show crypto ssh-key

Display the current host public keys used in SSH authentication.

Syntax `show crypto ssh-key {rsa | ecdsa | ed25119}`

Parameters

- `rsa` — Displays the RSA public key.
- `ecdsa` — Displays the ECDSA public key.
- `ed25519` — Displays the Ed25519 key.

Default Not configured

Command Mode EXEC

Usage Information After you regenerate an SSH server key with a customized bit size, disable and re-enable the SSH server to use the new publi keys. Use the `show crypto` command to verify the changes.

If a remote SSH client uses strict host-key checking, copy a newly generated host key to the list of known hosts on the client device.

Example

```
OS10# show crypto ssh-key rsa
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQCogJtArA0fHJkFpioGaAcp+vrDQFC3l3XFHtd41wXY9kM0Ar
+37yRsDul8vKodqSDiGLRuPjFTcVjvDdSKWblJRsybkmA6nuHJIyPOScDepLlicMIOxDhXEE92VRAmGuLI2AoeV
+IneWXhwQOkOFLtpxfnsiQY65CfS4aGoHOHWSfX3wI7boEDRDuvZ8gzRxTuM16Qr+RxBLJ7/OzkjNIN1/8Ok
+8aJtCoJKbcYaduMjmhVNrNUW5TUXoCnp1XNRpkJzgS7Lt47yi86rqrTCAQW4eSYJIJs4
+4ql9b4MF2D3499Ofn8uS82Mjtj0Nl01lbTbP3gsF4YYdBWaFqp root@OS10
```

Supported Releases 10.4.1.0 or later

show ip ssh

Displays the SSH server information.

Syntax `show ip ssh`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information Use this command to view information about the established SSH sessions.

Example

```
OS10# show ip ssh

SSH Server:                     Enabled
-------------------------------------------------
SSH Server Ciphers:             chacha20-poly1305@openssh.com,aes128-ctr,
                                aes192-ctr,aes256-ctr,
                                aes128-gcm@openssh.com,aes256-gcm@openssh.com
SSH Server MACs:                umac-64-etm@openssh.com,umac-128-etm@openssh.com,
                                hmac-sha2-256-etm@openssh.com,
                                hmac-sha2-512-etm@openssh.com,
                                hmac-sha1-etm@openssh.com,umac-64@openssh.com,
                                umac-128@openssh.com,hmac-sha2-256,
                                hmac-sha2-512,hmac-sha1
SSH Server KEX algorithms:      curve25519-sha256@libssh.org,ecdh-sha2-nistp256,
                                ecdh-sha2-nistp384,ecdh-sha2-nistp521,
                                diffie-hellman-group-exchange-sha256,
                                diffie-hellman-group14-sha1
Password Authentication:        Enabled
Host-Based Authentication:      Disabled
RSA Authentication:             Enabled
Challenge Response Auth:        Disabled
```

Supported Releases 10.3.0E or later

show logging audit

Displays audit log entries.

Syntax `show logging audit [reverse] [number]`

Parameters

- `reverse` — Display entries starting with the most recent events.
- *number* — Display specified number of audit log entries users, from 1 to 65535.

Default Display 24 entries starting with the oldest events.

Command Mode EXEC

Usage Information Only the `sysadmin` and `secadmin` roles can display the audit log. Enter `reverse` to display entries starting with the most recent events. You can change the number of entries displayed. Audit log records are not displayed on the console as they occur. They are saved in the audit log and forwarded to any configured Syslog servers.

Example

```
OS10# show logging audit 4
<14>1 2019-02-14T13:15:06.283337+00:00 OS10 audispd - - - Node.1-Unit.1:PRI [audit], Dell EMC
(OS10)  node=OS10 type=USER_END msg=audit(1550150106.277:597): pid=7908 uid=0 auid=4294967295
ses=4294967295 msg='op=PAM:session_close acct="admin" exe="/bin/su" hostname=? addr=?
```

```
terminal=??? res=success'
<110>1 2019-02-14T13:15:16.331515+00:00 OS10 .clish 7412 - -  Node.1-Unit.1:PRI [audit], User
admin on console used cmd: 'crypto security-profile mltestprofile' - success
<110>1 2019-02-14T13:15:21.794529+00:00 OS10 .clish 7412 - -  Node.1-Unit.1:PRI [audit], User
admin on console used cmd: 'exit' - success
<110>1 2019-02-14T13:16:05.882555+00:00 OS10 .clish 7412 - -  Node.1-Unit.1:PRI [audit], User
admin on console used cmd: 'exit' - success

OS10# show logging audit reverse 4
<110>1 2019-02-14T13:16:05.882555+00:00 OS10 .clish 7412 - -  Node.1-Unit.1:PRI [audit], User
admin on console used cmd: 'exit' - success
<110>1 2019-02-14T13:15:21.794529+00:00 OS10 .clish 7412 - -  Node.1-Unit.1:PRI [audit], User
admin on console used cmd: 'exit' - success
<110>1 2019-02-14T13:15:16.331515+00:00 OS10 .clish 7412 - -  Node.1-Unit.1:PRI [audit], User
admin on console used cmd: 'crypto security-profile mltestprofile' - success
<14>1 2019-02-14T13:15:06.283337+00:00 OS10 audispd - - - Node.1-Unit.1:PRI [audit], Dell EMC
(OS10)  node=OS10 type=USER_END msg=audit(1550150106.277:597): pid=7908 uid=0 auid=4294967295
ses=4294967295 msg='op=PAM:session_close acct="admin" exe="/bin/su" hostname=? addr=?
terminal=??? res=success'OS10# show logging audit reverse 10
```

Supported Releases

10.4.3.0 or later

show login-statistics

Displays statistics on user logins to the system.

Syntax `show login-statistics {user user-id | all}`

Parameters

- `user user-id` — Enter an OS10 username.
- `all` — Displays login statistics for all system users.

Default Not configured

Command Mode EXEC

Usage Information Only the `sysadmin` and `secadmin` roles can access this command. The `show` output displays login information for system users, including the number of successful and failed logins, role changes, and the last time a user logged in.

Example

```
OS10# show login-statistics all
Display statistics upon user login: Enabled
Time-frame in days: 25

               #Fail
               since During
        Role   last  Timeframe                Last Login
User    Change Login #Fail #Success  Date/Time             Location
-------- ----- ----- --------------  ------------------    ----------
admin    False  0     1    13        2017-11-02T16:02:44Z  in
netadmin False  0     0     5        2017-11-02T15:59:04Z  (00:00)
mltest   False  0     0     1        2017-11-01T15:42:07Z  1001:10:16:210::4001


OS10# show login-statistics user mltest
User                          : mltest
Role changed since last login : False
Failures since last login     : 0
Time-frame in days            : 25
Failures in time period       : 0
Successes in time period      : 1
```

```
Last Login Time                : 2017-11-01T15:42:07Z
Last Login Location            : 1001:10:16:210::4001
```

Supported Releases

10.4.0E(R1) or later

show privilege

Displays your current privilege level.

Syntax	`show privilege`
Parameters	None
Defaults	Not configured
Command Mode	EXEC
Example	`OS10# show privilege` `Current privilege level is 15.`
Supported Releases	10.4.3.0 or later

show running-configuration privilege

Displays the configured privilege levels of all users.

Syntax	`show running-configuration privilege`
Parameters	None
Defaults	Not configured
Command Mode	EXEC
Example	`OS10# show running-configuration privilege`
Supported Releases	10.4.3.0 or later

show users

Displays information for all users logged into OS10.

Syntax	`show users`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	None

Example	`OS10# show users`
	
	`Index Line   User  Role     Application Idle Login-Time            Location         Privilege`
	`----- ------ ----- -----   ----------- ---- ------------            --------         ---------`
	`1     pts/0  admin sysadmin bash        >24h 2018-09-08 T06:51:37Z 10.14.1.91 [ssh] 15`
	`2     pts/1  netad netadmin bash        >24h 2018-09-08 T06:54:33Z 10.14.1.91 [ssh] 10`

Supported Releases	10.2.0E or laterUpdated the command to display the privilege levels of all users on OS10 version 10.4.3.0 or later.

system-user linuxadmin disable

Enable or Disable and Lock or Unlock the linuxadmin user.

Syntax

```
[no] system-user linuxadmin disable
```

Parameters

None

Defaults

Enabled

Command Mode

CONFIGURATION

Usage Information

Use this command to disable and lock the linuxadmin user. The no version of this command enables or unlocks the linuxadmin user.

Example

```
OS10(config)# system-user linuxadmin disable

OS10(config)# no system-user linuxadmin disable
```

Supported Releases

10.4.3.0 or later

system-user linuxadmin password

Configure a password for the linuxadmin user.

Syntax

```
system-user linuxadmin password {clear-text-password | hashed-password}
```

Parameters

None

Defaults

Not configured

Command Mode

CONFIGURATION

Usage Information

Set a clear-text or hashed-password for the linuxadmin user.

Example

```
OS10(config)# system-user linuxadmin password Dell@Force10!@

OS10(config)# system-user linuxadmin password $6$3M55wOYy
$Sw1V9Ok3GE4Hmf6h1ARH.dBHy9gpEFYUvdu15ZpnCYzt.nJjFm0VIz/
rQvvJeX6krRtfYs2ZqBl6TkmLGAwtM
```

Supported Releases

10.4.3.0 or later

tacacs-server host

Configures a TACACS+ server and the key used to authenticate the switch on the server.

Syntax

```
tacacs-server host {hostname | ip-address} key {0 authentication-key | 9
authentication-key | authentication-key} [auth-port port-number]
```

Parameters

- *hostname* — Enter the host name of the TACACS+ server.
- *ip-address* — Enter the IPv4 (A.B.C.D) or IPv6 (x:x:x:x::x) address of the TACACS+ server.
- `key 0` *authentication-key* — Enter an authentication key in plain text. A maximum of 42 characters.
- `key 9` *authentication-key* — Enter an authentication key in encrypted format. . A maximum of 128 characters.

- *authentication-key* — Enter an authentication in plain text. . A maximum of 42 characters. It is not necessary to enter 0 before the key.
- key *authentication-key* — Enter a text string for the encryption key used to authenticate the switch on the TACACS+ server. A maximum of 42 characters.

Default Not configured

Command Mode CONFIGURATION

Usage Information The authentication key must match the key configured on the TACACS+ server. You cannot enter spaces in the key. The `show running-configuration` output displays both unencrypted and encrypted keys in encrypted format. Configure the global timeout allowed for authentication requests on TACACS+ servers using the `tacacs-server timeout` command. By default, OS10 times out an authentication attempt on a TACACS+ server after five seconds. The `no` version of this command removes a TACACS+ server configuration.

Example

```
OS10(config)# tacacs-server host 1.5.6.4 key secret1
```

Supported Releases 10.4.0E(R2) or later

tacacs-server timeout

Configures the global timeout used for authentication attempts on TACACS+ servers.

Syntax `tacacs-server timeout` *seconds*

Parameters *seconds* — Enter the timeout period used to wait for an authentication response from a TACACS+ server, from 1 to 1000 seconds.

Default 5 seconds

Command Mode CONFIGURATION

Usage Information The `no` version of this command resets the TACACS+ server timeout to the default.

Example

```
OS10(config)# tacacs-server timeout 360
```

Supported Releases 10.4.0E(R2) or later

username password role

Creates an authentication entry based on a user name and password, and assigns a role to the user.

Syntax `username` *username* `password` *password* `role` *role* `[ priv-lvl` *privilege-level*`]`

Parameters

- `username` *username*—Enter a text string. A maximum of 32 alphanumeric characters; one character minimum.
- `password` *password*—Enter a text string. A maximum of 32 alphanumeric characters; nine characters minimum.
- `role` *role*—Enter a user role:
  - `sysadmin`—Full access to all commands in the system, exclusive access to commands that manipulate the file system, and access to the system shell. A system administrator can create user IDs and user roles.
  - `secadmin`—Full access to configuration commands that set security policy and system access, such as password strength, AAA authorization, and cryptographic keys. A security administrator can display security information, such as cryptographic keys, login statistics, and log information.

  - netadmin—Full access to configuration commands that manage traffic flowing through the switch, such as routes, interfaces, and ACLs. A network administrator cannot access configuration commands for security features or view security information.
  - netoperator—Access to EXEC mode to view the current configuration. A network operator cannot modify any configuration setting on a switch.
- priv-lvl *privilege-level*—Enter a privilege level, from 0 to 15. If you do not specify the priv-lvl option, the system assigns privilege level 1 for the netoperator user and privilege level 15 for the sysadmin, secadmin, and netadmin users.

Default

- User name and password entries are in clear text.
- There is no default user role.

Command Mode CONFIGURATION

Usage Information

- By default, the password must be at least nine alphanumeric characters. Only the following special characters are supported:

  ```
  ! # % & ' ( ) ; < = > [ ] * + - . / : ^ _
  ```

  Enter the password in clear text. It is converted to SHA-512 format in the running configuration. For backward compatibility with OS10 releases 10.3.1E and earlier, passwords entered in MD-5, SHA-256, and SHA-512 format are supported.
- To increase the required password strength, use the password-attributes command.
- You cannot assign a privilege level higher than privilege level 1 to a user with the netoperator role and below and privilege level 2 for a sysadmin, secadmin, and netadmin roles.
- The no version of this command deletes authentication for a user.

Example

```
OS10(config)# username user05 password newpwd404 role sysadmin
OS10(config)# username user07 password paSsw0$d role netadmin priv-lvl 10
```

Supported Releases 10.2.0E or laterIntroduced the priv-lvl command on OS10 release 10.4.3.0 or later

username sshkey

Enables SSH password-less login using the public key for a remote client. The remote client is not prompted to enter a password.

Syntax `username *user_name* sshkey *sshkey_string*`

Parameters

- *user_name* — Enter the user name of the remote client. This value is the user name configured with the username password role command.
- *sshkey_string* — Enter the public key used by the remote client device to log in to the OS10 switch.

Default The default SSH server keys are an RSA key generated using 2048 bits, an ECDSA key with 256 bits, and an Ed2559 key with 256 bits.

Command Mode CONFIGURATION

Usage Information Locate the public keys on a remote client in the ~/.ssh/id_rsa.pub file. Use the public key as the *sshkey_string* parameter.

The no username *user_name* sshkey command removes the SSH password-less configuration for a specified user name.

To configure multiple user names for SSH password-less login, use the `username sshkey filename` command.

Example

```
OS10(config)# username user10 sshkey abcd

OS10(config)# do show running-configuration users
username admin password $6$q9QBeYjZ$jfxzVqGhkxX3smxJSH9DDz7/3OJc6m5wjF8nnLD7/
VKx8SloIhp4NoGZs0I/UNwh8WVuxwfd9q4pWIgNs5BKH. role sysadmin
username user10 password $6$rounds=656000$G10VRFTJB291ekwo
$iTGf0zd4bTUcBBpIVsbr6oStnUZMydN5lDs4WE6G3XHEtWbcKrGTeAo1wEF0cenEgRRPzi3SMmYyzAHCCC8wS0
role sysadmin
username user10 sshkey abcd
```

Supported Releases 10.4.1.0 or later

username sshkey filename

Enables SSH password-less login for remote clients using multiple public keys. A remote client is not prompted to enter a password.

Syntax `username user_name sshkey filename file_path`

Parameters

- *user_name* — Enter an OS10 user name who logs in on a remote client. This value is the user name configured with the `username password role` command.
- *file_path* — Enter the absolute path name of the local file containing the public keys used by remote devices to log in to the OS10 switch.

Default The default SSH server keys are an RSA key generated using 2048 bits, an ECDSA key with 256 bits, and an Ed2559 key with 256 bits.

Command Mode CONFIGURATION

Usage Information Before you use the command, locate the public keys on a remote client in the `~/.ssh/id_rsa.pub` file. Create a text file and copy the SSH public keys on the remote client into the file. Enter each public key on a separate line. Download the file to your home OS10 directory.

NOTE: Entering the command when an SSH key file is not present has no effect and results in a silent failure. SSH password-less login is not enabled.

The `no username user_name sshkey` command removes the SSH password-less configuration for the specified user name.

Example

```
OS10(config)# username user10 sshkey filename /test_file.txt

OS10(config)# do show running-configuration users
username admin password $6$q9QBeYjZ$jfxzVqGhkxX3smxJSH9DDz7/3OJc6m5wjF8nnLD7/
VKx8SloIhp4NoGZs0I/UNwh8WVuxwfd9q4pWIgNs5BKH. role sysadmin
username user10 password $6$rounds=656000$G10VRFTJB291ekwo
$iTGf0zd4bTUcBBpIVsbr6oStnUZMydN5lDs4WE6G3XHEtWbcKrGTeAo1wEF0cenEgRRPzi3SMmYyzAHCCC8wS0
role sysadmin
username user10 sshkey filename /test_file.txt
```

Supported Releases 10.4.1.0 or later

userrole inherit

Reconfigures the default `netoperator` role and permissions that OS10 assigns by default to a RADIUS or TACACS+-authenticated user with an unknown user role or privilege level. You can also configure an unknown RADIUS or TACACS+ user role to inherit permissions from an existing OS10 role.

Syntax

```
userrole {default | name} inherit existing-role-name
```

Parameters

- `default inherit` — Reconfigure the default permissions assigned to an authenticated user with a missing or unknown role or privilege level.
- *`name`* `inherit` — Enter the name of the RADIUS or TACACS+ user role that inherits permissions from an OS10 user role; 32 characters maximum.
- *`existing-role-name`* — Assign the permissions associated with an OS10 user role:
  - `sysadmin` — Full access to all commands in the system, exclusive access to commands that manipulate the file system, and access to the system shell. A system administrator can create user IDs and user roles.
  - `secadmin` — Full access to configuration commands that set security policy and system access, such as password strength, AAA authorization, and cryptographic keys. A security administrator can display security information, such as cryptographic keys, login statistics, and log information.
  - `netadmin` — Full access to configuration commands that manage traffic flowing through the switch, such as routes, interfaces, and ACLs. A network administrator cannot access configuration commands for security features or view security information.
  - `netoperator` — Access only to EXEC mode to view the current configuration. A network operator cannot modify any configuration setting on a switch.

Default OS10 assigns the `netoperator` role to a user authenticated by a RADIUS or TACACS+ server with a missing or unknown role or privilege level.

Command Mode CONFIGURATION

Usage Information

- When a RADIUS or TACACS+ server authenticates a user and does not return a role or privilege level, or returns an unknown role or privilege level, OS10 assigns the `netoperator` role to the user by default. Use this command to reconfigure the default `netoperator` permissions.
- To assign OS10 user role permissions to an unknown user role, enter the RADIUS or TACACS+ *name* with the `inherit` *`existing-role-name`* value. The `no userrole default` version of the command resets the role to `netoperator`.

Example

```
OS10(config)# userrole default inherit sysadmin
```

Supported Releases 10.4.0E(R3P3) or later

X.509v3 certificates

OS10 supports X.509v3 certificates to secure communications between the switch and a host, such as a RADIUS server. Both the switch and the server exchange a public key in a signed X.509v3 certificate issued by a certificate authority (CA) to authenticate each other. The certificate authority uses its private key to sign the switch and host certificates.

The information in the certificate allows both devices to prove ownership and the validity of a public key. Assuming the CA is trusted, the switch and authentication server validate each other's identity and set up a secure, encrypted communications channel.

User authentication with a public key certificate is usually preferred to password-based authentication, although you can use both at the same time, to:

- Avoid the security risk of using low-strength passwords and provide greater resistance to brute-force attacks.
- Provide assurance of trusted, provable identities (when using certificates digitally signed by a trusted CA).
- Provide security and confidentiality in switch-server communications in addition to user authentication.

For example, you can download and install a X.509v3 certificate to enable public-key authentication in RADIUS over TLS authentication, also called as `RadSec`. OS10 supports a public key infrastructure (PKI), including:

- Generation of self-signed certificates and certificate signing requests (CSRs), and their corresponding private keys
- Installation and deletion of self-signed certificates and CA-signed certificates
- Secure deletion of corresponding private keys
- Installation and deletion of CA certificates in the system "trust store"
- Display of certificate information

X.509v3 concepts

Certificate	A document that associates a network device with its public key. When exchanged between participating devices, certificates are used to validate device identity and the public key associated with the device. A PKI uses the following certificate types: • CA certificate: The certificate of a CA that is used to sign host certificates. A CA certificate may be issued by other CAs or be self-signed. A self-signed CA certificate is called a *root certificate*. • Host certificate: A certificate that is issued to a network device. A host certificate can be signed by a CA or self-signed. • Self-signed certificate: A host-signed certificate, compared to a CA-signed certificate.
Certificate authority (CA)	An entity that verifies the contents of a certificate and signs it, indicating that the certificate is trusted and correct. An intermediate CA signs certificates transmitted between a root CA and a host.
Certificate revocation list (CRL)	A CA-signed document that contains a list of certificates that are no longer valid, even though they have not yet expired. For example, when a new certificate is generated for a server, and the old certificate is no longer supported.
Certificate signing request (CSR)	After generating a key pair, a switch signs a request to obtain a certificate using its secret private key, and sends the request to a certificate authority. The CSR contains information that identifies the switch and its public key. This public key is used to verify the private signature of the CSR and the distinguished name (DN) of the switch. A CSR is signed by a CA and returned to a host for use as a signed host certificate.
Privacy Enhanced Mail (PEM)	PKI standard used to format X.509v3 data in a secure message exchange; described in RFC 1421.
Public key infrastructure (PKI)	Application that manages the generation of private and public encryption keys, and the download, installation, and exchange of CA-signed certificates with network devices.
X.509v3	Standard for the public key infrastructure that manages digital certificates and public key encryption.

Public key infrastructure

To use X.509v3 certificates for secure communication and user authentication on OS10 switches in a network, a public key infrastructure (PKI) with a certificate authority (CA) is required. The CA signs certificates that prove the trustworthiness of network devices.

When an organization wants to assure customers that the connection to their network is secure, it may pay a commercial Certificate Authority, such as VeriSign or DigiCert, to sign a certificate for their domain. However, to implement an X.509v3 infrastructure, you can act as your own CA. While acting as your own CA, you can set up CAs to issue certificates to hosts in the same trusted domain to authenticate each other.

X.509v3 public key infrastructure

To set up a PKI using X.509v3 certificates, Dell EMC Networking recommends:

1 Configure a root CA that generates a private key and a self-signed CA certificate.

2 Configure one or more intermediate CAs that generate a private key and a certificate signing request (CSR), and send the CSR to the root CA.
   - Using its private key, the root CA signs an intermediate CA's CSR and generates a CA certificate for the Intermediate CA.
   - The intermediate CA downloads and installs the CA certificate. Afterwards, the intermediate CA can sign certificates for hosts in the network and for other intermediate CAs that are lower in the PKI hierarchy.
   - The root and intermediate CA certificates, but not the corresponding private keys, are made publicly available on the network for network hosts to download.
   - Whenever possible, store private keys offline or in a location restricted from general access.

3 OS10 switches generate private keys and create CSRs using the `crypto cert generate request` command. A switch uploads a CSR to an intermediate CA. To store the private key in a local hidden location, Dell EMC Networking recommends using the `key-file private` parameter with the command.

4 Download and install a CA certificate on a host using the `crypto ca-cert install` command. After you install a CA certificate, a host trusts any certificates that are signed by the CA and presented by other network devices. You must first download a certificate to the home directory, and then install the certificate using the `crypto ca-cert install` command.

5 Download and install a signed host certificate and private key from an intermediate CA on an OS10 switch. Then install them using the `crypto cert install` command. After you install the host certificate, OS10 applications use the certificate to secure communication with network devices. The private key is installed in the internal file system on the switch and cannot be exported or viewed.

Manage CA certificates

OS10 supports the download and installation of public X.509v3 certificates from external certificate authorities.

In a data center environment, trusted CA servers can create CA certificates. A host operates as a trusted CA server. Network hosts install certificates that are digitally signed with the CA's private key to establish trust between participating devices in the network. The certificate on an OS10 switch is used to verify the certificates presented by clients and servers, such as Syslog and RADIUS servers, to establish a secure connection with these devices.

To import a CA server certificate:

1 Use the `copy` command to download an X.509v3 certificate created by a CA server using a secure method, such as HTTPS, SCP, or SFTP. Copy a CA certificate to the local directory on the switch, such as `home://` or `usb://`.

2 Use the `crypto ca-cert install` command to install the certificate. When you install a CA certificate, specify the local path where the certificate is stored.

The switch verifies the certificate and installs it in an existing directory of trusted certificates in PEM format.

Install CA certificate

- Install a CA certificate in EXEC mode.

  ```
  crypto ca-cert install ca-cert-filepath [filename]
  ```

  - *ca-cert-filepath* specifies the local path to the downloaded certificate; for example, `home://CAcert.pem` or `usb://CA-cert.pem`.
  - *filename* specifies an optional filename that the certificate is stored under in the OS10 trust-store directory. Enter the filename in the *filename*`.crt` format.

Example: Download and install CA certificate

```
OS10# copy scp:///tftpuser@10.11.178.103:/tftpboot/certs/Dell_rootCA1.pem home://
Dell_rootCA1.pem
password:

OS10# crypto ca-cert install home://Dell_rootCA1.pem
Processing certificate ...
Installed Root CA certificate
  CommonName = Dell_rootCA1
  IssuerName = Dell_rootCA1
```

Display CA server certificate

```
OS10# show crypto ca-certs
 -------------------------------------
|    Locally installed certificates    |
 -------------------------------------
Dell_rootCA1.crt
```

```
OS10# show crypto ca-certs Dell_rootCA1.crt
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number:
            95:48:23:17:76:9d:05:e1
    Signature Algorithm: sha256WithRSAEncryption
        Issuer: C = US, ST = California, L = Santa Clara, O = Dell EMC, OU = Networking, CN =
Dell_rootCA1
        Validity
            Not Before: Jul 25 18:21:50 2018 GMT
            Not After : Jul 20 18:21:50 2038 GMT
        Subject: C = US, ST = California, L = Santa Clara, O = Dell EMC, OU = Networking, CN =
Dell_rootCA1
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (4096 bit)
                Modulus:
                    00:cd:9d:ca:10:6b:b1:54:81:10:92:42:9f:6a:cb:
                    49:51:9d:46:10:cb:67:08:2b:75:2a:62:40:80:a3:
                    f5:7d:58:67:f4:cc:c6:70:32:14:4c:f0:4d:cd:7e:
                    0d:5c:63:28:5e:6c:ad:9e:13:13:71:6d:9d:10:a9:
                    a1:d8:6b:bd:a3:a0:5a:11:19:87:4d:3d:08:6f:10:
                    03:df:70:89:5f:b7:56:49:32:57:9c:28:5e:43:7f:
                    ca:bc:41:c7:31:51:97:7f:73:b7:b0:c4:13:21:e6:
                    2c:4c:19:fd:35:0b:26:16:78:fc:c3:73:21:3a:06:
                    f6:ec:87:3f:9f:5e:3a:0c:23:5e:13:4c:9e:5a:70:
                    18:d4:ad:cb:cf:47:c1:c6:50:a0:49:df:a0:a6:47:
                    1e:13:19:49:9e:67:db:1c:c7:23:9e:37:3b:c7:0c:
                    cd:26:46:f6:c1:e1:93:64:29:81:9c:e9:a8:1d:29:
                    19:4c:8d:a4:a8:53:66:2b:b2:70:ff:ec:80:d4:87:
                    eb:74:e2:11:56:ed:4b:68:fc:53:2e:d4:94:f6:f5:
                    e4:77:d9:b6:e8:4a:91:b7:da:46:18:51:bf:e4:b6:
                    3e:6a:47:ab:77:f6:93:b7:d0:9a:c8:fa:ba:ae:ed:
                    6a:fd:81:54:c8:76:13:1b:57:74:d6:02:78:d7:98:
                    38:e6:c5:9b:64:03:b2:76:93:fd:8c:9f:54:c9:a3:
                    04:a9:0c:b7:e2:bb:02:50:3f:e0:08:33:32:89:55:
                    95:9b:30:6c:73:7d:be:63:f1:6c:da:4d:92:41:d0:
                    f5:d6:bf:e3:c0:da:98:ae:24:37:ed:07:63:86:a1:
                    cc:da:3b:45:d4:a9:80:e2:d6:ab:c1:ae:2a:99:32:
                    9d:ba:fe:88:38:f2:02:d1:b3:78:43:17:7e:6e:b1:
                    a2:17:85:bd:5f:4a:52:90:96:4d:bc:19:85:ed:9d:
                    49:77:bd:59:44:6c:6c:23:e5:b1:92:af:a0:10:ce:
                    68:d4:f4:07:9e:ec:ca:c5:95:a2:f4:19:bb:f7:12:
                    ce:f0:a6:39:df:1a:5b:10:91:d5:77:46:8d:55:9a:
                    8e:96:e0:70:f6:27:89:43:3d:74:99:b4:7f:4b:38:
                    71:18:01:64:bb:72:2c:26:6f:6e:e8:06:9a:77:4b:
                    07:3b:b3:8c:71:ff:61:1b:84:d4:02:46:47:e5:4d:
                    79:be:22:e9:7a:8c:eb:06:38:38:a6:f7:b7:83:bf:
                    f2:64:c9:b8:d9:7f:d1:cc:87:ac:80:b0:d0:d3:17:
                    35:d1:49:44:2e:6e:9f:60:9c:ca:9a:6d:cd:63:79:
                    7c:6d:33:72:13:74:f1:16:20:50:46:20:e7:c1:ff:
                    b0:42:95
                Exponent: 65537 (0x10001)
        X509v3 extensions:
            X509v3 Subject Key Identifier:
                75:22:3F:BE:99:B7:FA:A1:5B:1D:68:0B:E9:5E:21:7D:83:62:AC:DB
            X509v3 Authority Key Identifier:
                keyid:75:22:3F:BE:99:B7:FA:A1:5B:1D:68:0B:E9:5E:21:7D:83:62:AC:DB
            X509v3 Basic Constraints: critical
                CA:TRUE
```

```
            X509v3 Key Usage: critical
                Digital Signature, Certificate Sign, CRL Sign
    Signature Algorithm: sha256WithRSAEncryption
         8e:0c:50:18:5f:db:cc:80:5c:6e:ce:43:29:32:2e:0b:70:96:
         db:e8:23:c9:15:a2:99:72:d6:01:c9:61:8e:ed:8d:f8:4d:2f:
         99:57:bf:52:1f:4a:5b:7b:ff:24:23:5f:eb:3e:e8:8e:0c:d4:
         94:0f:20:a7:e3:3b:18:e9:76:06:5a:ae:65:38:d4:3a:98:d6:
         0b:73:5b:b5:8e:4c:b5:74:02:9a:9d:9a:7d:7a:18:2f:32:38:
         9e:0e:7b:de:15:3c:f1:33:e8:2d:3f:92:f0:f2:4e:7a:7f:e2:
         a5:2e:04:3a:2f:3b:1b:05:71:39:70:6d:a4:6e:8f:25:31:0e:
         2c:8a:7e:b4:30:7c:38:2f:48:df:19:56:42:4f:be:5f:d3:02:
         70:18:7e:76:66:ca:13:1c:e3:9c:4d:aa:d3:67:96:be:d9:49:
         5c:69:10:75:26:53:f7:50:39:06:15:d1:3a:87:47:f6:92:a2:
         d4:91:35:29:b7:4b:ea:56:4c:13:5e:32:7f:c7:3f:4c:46:67:
         54:8d:67:60:38:98:75:da:24:f2:64:b9:24:a1:e3:5b:42:66:
         4c:c7:cb:ee:c3:ca:bd:87:1b:7a:fc:35:53:2d:74:68:db:a7:
         47:db:03:a3:30:52:af:67:7f:54:a4:de:60:ca:ae:94:43:f8:
         98:85:fc:18:9b:b1:db:81:44:57:0b:be:6a:56:9d:2f:7d:75:
         c2:22:a4:7c:d7:ee:f8:de:10:11:26:60:35:1c:4c:87:2e:a2:
         fb:1f:5f:30:6c:11:c1:fa:f2:5b:46:02:0a:18:2f:02:a4:99:
         f2:43:29:cf:e6:5b:8a:d0:ec:42:bf:49:c6:8a:7e:b4:53:38:
         03:1b:fd:a9:49:88:b5:f1:42:93:c7:78:38:6c:2a:1c:be:83:
         97:27:b1:26:eb:16:44:ce:34:02:53:45:08:30:c9:3a:76:83:
         10:f3:af:c7:6f:0c:74:ec:81:ea:d9:c4:20:a5:1d:72:64:52:
         7b:e8:30:1a:9e:3a:05:9c:8a:69:e5:b7:43:b3:36:08:f2:e0:
         fb:88:d9:c1:b6:f4:4a:23:27:31:3a:51:b3:68:c9:6f:3e:f5:
         dd:98:4d:07:38:ed:f4:d3:ed:06:4c:84:87:3d:cf:f3:2e:e5:
         1a:b6:00:71:4c:51:35:c8:95:e4:c6:7e:82:47:d3:25:64:a4:
         0b:31:53:d0:e4:6b:97:98:21:4b:fc:e7:12:be:69:01:d8:b5:
         74:f5:b6:39:22:8a:8c:39:23:0f:be:4b:0f:9a:01:ac:b8:5b:
         12:cb:94:06:30:f5:74:45:20:af:ab:d6:af:21:0c:d8:62:84:
         18:c2:cf:4f:be:73:c9:33
```

Delete CA server certificate

```
OS10# crypto ca-cert delete Dell_rootCA1.crt
Successfully removed certificate
```

Request and install host certificates

OS10 also supports the switch obtaining its own X.509v3 host certificate. In this procedure, you generate a certificate signing request (CSR) and a private key. Store the private key locally in a secure location. Copy the CSR file to a certificate authority. The CA generates a host certificate for an OS10 switch by digitally signing the switch certificate contained in the CSR.

The administrator then copies the CA-signed host certificate to the home directory on the switch. Because a local private key is created when the CSR is generated, it is not necessary to install a private key using an uploaded file.

The switch presents its own host certificate to clients that require authentication, such as Syslog and RADIUS servers over TLS and HTTPS connections. The certificate is digitally signed with the private key of the OS10 switch. OS10 supports multiple host certificates so that you can use different certificates with different applications. For more information, see Security profiles.

To obtain a host certificate from a CA:

1 Create a private key and generate a certificate signing request for the switch.
2 Copy the CSR file to a CA server.
3 Copy the CA-signed certificate to the home directory on the switch. Install the trusted certificate.

Generate a certificate signing request and private key

- Create a private key and a CSR in EXEC mode. Store the CSR file in the home directory or `flash:` so that you can later copy it to a CA server. Specify a *keypath* to store the `device.key` file in a secure persistent location, such as the home directory, or use the `private` option to store the key file in a private hidden location in the internal file system that is not visible to users.

  ```
  crypto cert generate request [cert-file cert-path key-file {private | keypath}]
  [country 2-letter code] [state state] [locality city] [organization organization-name]
  [orgunit unit-name] [cname common-name] [email email-address] [validity days]
  [length length] [altname alt-name]
  ```

  If you enter the `cert-file` option, you must enter all the following required parameters, such as the local paths where the certificate and private key are stored, country code, state, locality, and other values.

  If you do not specify the `cert-file` option, you are prompted to fill in the other parameter values for the certificate interactively; for example:

  ```
  You are about to be asked to enter information that will be incorporated into your
  certificate request.
  What you are about to enter is what is called a Distinguished Name or a DN.
  There are quite a few fields but you can leave some blank.
  For some fields there will be a default value; if you enter '.', the field will be left blank.
  Country Name (2 letter code) [US]:
  State or Province Name (full name) [Some-State]:California
  Locality Name (eg, city) []:San Francisco
  Organization Name (eg, company) []:Starfleet Command
  Organizational Unit Name (eg, section) []:NCC-1701A
  Common Name (eg, YOUR name) [hostname]:S4148-001
  Email Address []:scotty@starfleet.com
  ```

  The switch uses SHA-256 as the digest algorithm. The public key algorithm is RSA with a 2048-bit modulus. The KeyUsage bits of the certificate assert `keyEncipherment` (bit 2) and `keyAgreement` (bit 4). The `keyCertSign` bit (bit 5) is NOT set. The ExtendedKeyUsage fields indicate `serverAuth` and `clientAuth`.

  The attribute `CA:FALSE` is set in the Extensions section of the certificate. The certificate is NOT used to validate other certificates.

- If necessary, re-enter the command to generate multiple certificate-key pairs for different applications on the switch. You can configure a certificate-key pair in a security profile. Using different certificate-key pairs is necessary if you want to change the certificate-key pair for a specified application without out interrupting other critical services. For example, RADIUS over TLS may use a different certificate-key pair than SmartFabric services.

NOTE:
If the system is in FIPS mode (`crypto fips enable` command), the CSR and private key are generated using FIPS-validated and compliant algorithms. You manage whether the keys are generated in FIPS mode or not.

Copy CSR to the CA server

You can copy the CSR from flash to a destination, such as a USB flash drive, using TFTP, FTP, or SCP.

```
OS10# copy home://DellHost.pem  scp:///tftpuser@10.11.178.103:/tftpboot/certs/DellHost.pem
password:
```

The CA server signs the CSR with its private key. The CA server then makes the signed certificate available for the OS10 switch to download and install.

Install host certificate

1. Use the `copy` command to download an X.509v3 certificate signed by a CA server to the local home directory using a secure method, such as HTTPS, SCP, or SFTP.
2. Use the `crypto cert install` command to install the certificate and the private key generated with the CSR.

- Install a trusted certificate and key file in EXEC mode.

  ```
  crypto cert install cert-file home://cert-filepath key-file {key-path | private}
  [password passphrase] [fips]
  ```

- `cert-file` *`cert-filepath`* specifies a source location for a downloaded certificate; for example, `home://s4048-001-cert.pem` or `usb://s4048-001-cert.pem`.
- `key-file {`*`key-path`*` | private}` specifies the local path to retrieve the downloaded or locally generated private key. Enter `private` to install the key from a local hidden location and rename the key file with the certificate name.
- `password` *`passphrase`* specifies the password used to decrypt the private key if it was generated using a password.
- `fips` installs the certificate-key pair as FIPS-compliant. Enter `fips` to install a certificate-key pair that is used by a FIPS-aware application, such as RADIUS over TLS. If you do not enter `fips`, the certificate-key pair is stored as a non-FIPS compliant pair.

  NOTE: You determine if the certificate-key pair is generated as FIPS-compliant. Make sure that FIPS-compliant certificate-key pairs are not used outside of FIPS mode. When FIPS mode is enabled on the switch, you can still generate CSRs for non-FIPS certificates for use with non-FIPS applications. Be sure to install these certificates as non-FIPS with the `crypto cert install` command.
- If you enter `fips` after using the `key-file private` option in the `crypto cert generate request` command, a FIPS-compliant private key is stored in a hidden location in the internal file system that is not visible to users.

If the certificate installation is successful, the file name of the host certificate and its common name are displayed. Use the filename to configure the certificate in a security profile (`crypto security-profile` command).

Example: Generate CSR and upload to server

```
OS10# crypto cert generate request cert-file home://DellHost.pem key-file home://DellHost.key
email admin@dell.com length 1024 altname DNS:dell.domain.com
Processing certificate ...
Successfully created CSR file /home/admin/DellHost.pem and key

OS10# copy home://DellHost.pem  scp:///tftpuser@10.11.178.103:/tftpboot/certs/DellHost.pem
password:
```

Host certificate tip

When administering a large number of switches, you may choose to not generate numerous CSRs for all switches. An alternate method to installing a host certificate on each switch is to generate both the private key file and CSR offline; for example, on the CA server. The CSR is signed by the CA, which generates both a certificate and key file. You then copy the trusted certificate and key file to the switch using the `copy` command and install them using the `crypto cert install cert-file home://`*`cert-filename`* `key-file home://`*`key-filename`* command.

NOTE: For security reasons, the private key file is copied to an internal, secure location and removed from the viewable file system.

Example: Download and install trusted certificate and private key

```
OS10# copy scp:///tftpuser@10.11.178.103:/tftpboot/certs/Dell_host1_CA1.pem home://
Dell_host1_CA1.pem
password:

OS10# copy scp:///tftpuser@10.11.178.103:/tftpboot/certs/Dell_host1_CA1.key home://
Dell_host1_CA1.key
password:

OS10# crypto cert install cert-file home://Dell_host1_CA1.pem key-file home://Dell_host1_CA1.key
Processing certificate ...
Certificate and keys were successfully installed as "Dell_host1_CA1.pem" that may be used in a
security profile. CN = Dell_host1_CA1
```

Display trusted certificates

```
OS10# show crypto cert
 ---------------------------------------
|    Installed non-FIPS certificates    |
 ---------------------------------------
Dell_host1_CA1.pem
 ---------------------------------------
```

```
|    Installed FIPS certificates    |
 ------------------------------------
```

```
OS10# show crypto cert Dell_host1_CA1.pem
 ------------ Non FIPS certificate -----------------
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 4096 (0x1000)
    Signature Algorithm: sha256WithRSAEncryption
        Issuer: C = US, ST = California, O = Dell EMC, OU = Networking, CN = Dell_interCA1
        Validity
            Not Before: Jul 25 19:11:19 2018 GMT
            Not After : Jul 22 19:11:19 2028 GMT
        Subject: C = US, ST = California, L = Santa Clara, O = Dell EMC, OU = Networking, CN =
Dell_host1_CA1
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (2048 bit)
                Modulus:
                    00:e7:81:4b:4a:12:8d:ce:88:e6:73:3f:da:19:03:
                    c6:56:01:19:b2:02:61:3f:5b:1e:33:28:a1:ed:e3:
                    85:bc:56:fb:18:d5:16:2e:a0:e7:3a:f9:34:b4:df:
                    37:97:93:a9:b9:94:b2:9f:69:af:fa:31:77:68:06:
                    89:7b:6d:fc:91:14:4a:c8:7b:23:93:f5:44:5a:0a:
                    3f:ce:9b:af:a6:9b:49:29:fd:fd:cb:34:40:c4:02:
                    30:95:37:28:50:d8:81:fb:1f:83:88:d9:1f:a3:0e:
                    49:a1:b3:df:90:15:d4:98:2b:b2:38:98:6e:04:aa:
                    bd:92:1b:98:48:4d:08:49:69:41:4e:6a:ee:63:d8:
                    2a:9f:e6:15:e2:1d:c3:89:f5:f0:d0:fb:c1:9c:46:
                    92:a9:37:b9:2f:a0:73:cf:e7:d1:88:96:b8:4a:84:
                    91:83:8c:f0:9a:e0:8c:6e:7a:fa:6e:7e:99:3a:c3:
                    2c:04:f9:06:8e:05:21:5f:aa:6e:9f:b7:10:37:29:
                    0c:03:14:a0:9d:73:1f:95:41:39:9b:96:30:9d:0a:
                    cb:d0:65:c3:59:23:01:f7:f5:3a:33:b9:e9:95:11:
                    0c:51:f4:e9:1e:a5:9d:f7:95:84:9c:25:74:0c:21:
                    4f:8b:07:29:2f:e3:47:14:50:8b:03:c1:fb:83:85:
                    dc:bb
                Exponent: 65537 (0x10001)
        X509v3 extensions:
            X509v3 Basic Constraints:
                CA:FALSE
            Netscape Cert Type:
                SSL Client, S/MIME
            Netscape Comment:
                OpenSSL Generated Client Certificate
            X509v3 Subject Key Identifier:
                4A:20:AA:E1:69:BF:BE:C5:66:2E:22:71:70:B4:7E:32:6F:E0:05:28
            X509v3 Authority Key Identifier:
                keyid:A3:39:CB:C7:76:86:3B:05:44:34:C2:6F:90:73:1F:5F:64:55:5C:76
            X509v3 Key Usage: critical
```

Delete trusted certificate

```
OS10# OS10# crypto cert delete Dell_host1_CA1.pem
Certificate and keys were successfully deleted. CN = Dell_host1_CA1
```

Self-signed certificates

Administrators may prefer to not set up a Certificate Authority and implement a certificate trust model in the network, but still want to use the privacy features provided by the Transport Layer Security (TLS) protocol. In this case, self-signed certificates can be used.

A self-signed certificate is not signed by a CA. The switch presents itself as a trusted device in its certificate. Connecting clients may prompt their users to trust the certificate — for example, when a web browser warns that a site is unsafe — or to reject the certificate, depending on the configuration. A self-signed certificate does not provide protection against man-in-the-middle attacks.

To generate and install a self-signed certificate:

1 Create a self-signed certificate and key in a local directory or USB flash drive.
2 Install the self-signed certificate.

Generate a self-signed certificate

- Create a self-signed certificate in EXEC mode. Store the `device.key` file in a secure, persistent location, such as NVRAM.

```
crypto cert generate self-signed [cert-file cert-path key-file {private | keypath}]
[country 2-letter code] [state state] [locality city] [organization organization-name]
[orgunit unit-name] [cname common-name] [email email-address] [validity days]
[length length] [altname alt-name]
```

If you enter the `cert-file` option, you must enter all the required parameters, including the local path where the certificate and private key are stored.

If you do specify the `cert-file` option, you are prompted to enter the other parameter values for the certificate interactively; for example:

```
You are about to be asked to enter information that will be incorporated in your certificate
request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank.
For some fields there will be a default value; if you enter '.', the field will be left blank.
Country Name (2 letter code) [US]:
State or Province Name (full name) [Some-State]:California
Locality Name (eg, city) []:San Francisco
Organization Name (eg, company) []:Starfleet Command
Organizational Unit Name (eg, section) []:NCC-1701A
Common Name (eg, YOUR name) [hostname]:S4148-001
Email Address []:scotty@starfleet.com
```

The switch uses SHA-256 as the digest algorithm. The public key algorithm is RSA with a 2048-bit modulus.

NOTE: When using self-signed X.509v3 certificates with Syslog and RADIUS servers, configure the server to accept self-signed certificates. Syslog and RADIUS servers require mutual authentication, which means that the client and server must verify each other's certificates. The best practice is to configure a CA server to sign certificates for all trusted devices in the network.

Install self-signed certificate

- Install a self-signed certificate and key file in EXEC mode.

```
crypto cert install cert-file home://cert-filename key-file {key-path | private}
[password passphrase] [fips]
```

  - `cert-file cert-path` specifies a source location for a downloaded certificate; for example, `home://s4048-001-cert.pem` or `usb://s4048-001-cert.pem`.
  - `key-file {key-path | private}` specifies the local path to retrieve the downloaded or locally generated private key. Enter `private` to install the key from a local hidden location and rename the key file with the certificate name.
  - `password passphrase` specifies the password used to decrypt the private key if it was generated using a password.
  - `fips` installs the certificate-key pair as FIPS-compliant. Enter `fips` to install a certificate-key pair that is used by a FIPS-aware application, such as RADIUS over TLS. If you do not enter `fips`, the certificate-key pair is stored as a non-FIPS compliant pair.

    NOTE: You determine if the certificate-key pair is generated as FIPS-compliant. Make sure that FIPS-compliant certificate-key pairs are not used outside of FIPS mode.
  - If you enter `fips` after using the `key-file private` option in the `crypto cert generate request` command, a FIPS-compliant private key is stored in a hidden location in the internal file system that is not visible to users.

If the certificate installation is successful, the file name of the self-signed certificate and its common name are displayed. Use the file name to configure the certificate in a security profile (`crypto security-profile` command).

Example: Generate and install self-signed certificate and key

```
OS10# crypto cert generate self-signed cert-file home://DellHost.pem key-file home://
DellHost.key email admin@dell.com length 1024 altname DNS:dell.domain.com validity 365
```

```
Processing certificate ...
Successfully created certificate file /home/admin/DellHost.pem and key

OS10# crypto cert install cert-file home://DellHost.pem key-file home://DellHost.key
Processing certificate ...
Certificate and keys were successfully installed as "DellHost.pem" that may be used in a
security profile. CN = DellHost.
```

Display self-signed certificate

```
OS10# show crypto cert
-------------------------------------
|    Installed non-FIPS certificates    |
-------------------------------------
DellHost.pem
-------------------------------------
|    Installed FIPS certificates    |
-------------------------------------
```

```
OS10# show crypto cert DellHost.pem
------------ Non FIPS certificate -----------------
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 245 (0xf5)
    Signature Algorithm: sha256WithRSAEncryption
        Issuer: emailAddress = admin@dell.com
        Validity
            Not Before: Feb 11 20:10:12 2019 GMT
            Not After : Feb 11 20:10:12 2020 GMT
        Subject: emailAddress = admin@dell.com
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (1024 bit)
                Modulus:
                    00:c7:12:ca:a8:d6:d2:1c:ab:66:9a:d1:db:50:5a:
                    b5:8a:e4:53:9d:f6:b4:fc:cd:f4:b9:46:8a:03:86:
                    be:0b:50:51:c7:25:76:9f:ff:b4:f9:f8:d9:6f:5d:
                    53:52:0c:4d:05:ed:31:23:79:44:5c:d7:62:01:9d:
                    41:e8:ff:3a:b0:35:0c:22:d7:ef:df:05:9a:28:6b:
                    95:10:8e:bc:c6:62:3a:82:30:0f:4f:4e:19:17:48:
                    f1:bd:1e:0c:4f:54:03:42:f3:a7:de:22:40:3d:5e:
                    6b:b2:8e:23:17:53:ef:10:d9:ae:1d:1f:d6:e4:ae:
                    25:9f:d9:39:60:5c:49:b0:ad
                Exponent: 65537 (0x10001)
        X509v3 extensions:
            X509v3 Subject Key Identifier:
                DA:39:A3:EE:5E:6B:4B:0D:32:55:BF:EF:95:60:18:90:AF:D8:07:09
            X509v3 Subject Alternative Name:
                DNS:dell.domain.com
    Signature Algorithm: sha256WithRSAEncryption
         b8:83:ae:34:bb:84:e6:b4:a3:fd:77:20:67:15:3f:02:76:ca:
         f6:74:d4:d2:36:0e:58:8c:96:13:c2:85:8a:df:ba:c0:d9:c8:
```

Security profiles

To use independent sets of security credentials for different OS10 applications, you can configure multiple security profiles and assign them to OS10 applications. A security profile consists of a certificate and private key pair.

For example, you can maintain different security profiles for RADIUS over TLS authentication and SmartFabric services. Using different security profiles allows you to upgrade one application without interrupting the operation of the other one. Assign a security profile to an application when you configure the profile.

When you install a certificate-key pair, both take the name of the certificate. For example, if you install a certificate using:

```
OS10# crypto cert install cert-file home://Dell_host1.pem key-file home://abcd.key
```

The certificate-key pair is installed as `Dell_host1.pem` and `Dell_host1.key`. In configuration commands, refer to the pair as `Dell_host1`. When you configure a security profile, you would enter `Dell_host1` in the `certificate` *`certificate-name`* command.

Configure security profile

1 Create an application-specific security profile in CONFIGURATION mode.

```
crypto security-profile profile-name
```

2 Assign a certificate and private key pair to the security profile in SECURITY-PROFILE mode. For *`certificate-name`*, enter the name of the certificate-key pair as it appears in the `show crypto certs` output without the `.pem` extension.

```
certificate certificate-name
exit
```

3 Use the security profile to configure X.509v3-based service; for example, to configure RADIUS over TLS authentication using an X.509v3 certificate, enter the `radius-server host tls` command:

```
radius-server host {hostname | ip-address} tls security-profile profile-name
[auth-port port-number] key {0 authentication-key | 9 authentication-key | authentication-key}
```

Example: Security profile configuration and use for RADIUS over TLS authentication

```
OS10# show crypto cert
--------------------------------------
|    Installed non-FIPS certificates    |
--------------------------------------
dv-fedgov-s6010-1.pem
--------------------------------------
|    Installed FIPS certificates    |
--------------------------------------
OS10#
OS10(config)#
OS10(config)# crypto security-profile radius-prof
OS10(config-sec-profile)# certificate dv-fedgov-s6010-1
OS10(config-sec-profile)# exit
OS10(config)#
OS10(config)# radius-server host radius-server-2.test.com tls security-profile radius-prof key
radsec
OS10(config)# end
OS10# show running-configuration crypto security-profile
!
crypto security-profile radius-prof
  certificate dv-fedgov-s6010-1

OS10# show running-configuration radius-server
radius-server host radius-server-2.test.com tls security-profile radius-prof key 9
2b9799adc767c0efe8987a694969b1384c541414ba18a44cd9b25fc00ff180e9
```

X.509v3 commands

certificate

Configures a certificate and private key pair in an application-specific security profile.

Syntax	`certificate` *`certificate-name`*
Parameters	*certificate-name* — Enter the name of the certificate-key pair as it appears in the `show crypto certs` output without the `.pem` extension.

Default Not configured

Command mode SEC-PROFILE

Usage information Use the `certificate` command to associate a certificate and private key with a security profile. An application-specific security profile allows you to change the certificate-key pair used by an OS10 application, such as SmartFabric services, without interrupting the service of other mission-critical applications.

When you install a certificate-key pair, both take the name of the certificate. Enter the certificate-key pair name without an extension as the *certificate-name* value. The `no` form of the command removes the certificate-key pair from the profile.

Example

```
OS10# crypto security-profile secure-radius-profile
OS10(config-sec-profile)# certificate Dell_host1
```

Supported releases 10.4.3.0 or later

crypto ca-cert delete

Deletes a CA certificate.

Syntax `crypto ca-cert delete {`*ca-cert-filepath* `| all}`

Parameters

- *ca-cert-filepath* — Enter the local path where the downloaded CA certificate is stored; for example, `home://CAcert.pem` or `usb://CA-cert.pem`.
- `all` — Delete all CA certificates.

Default Not configured

Command mode EXEC

Usage information To display the currently installed CA certificates, use the `show crypto ca-certs` command.

Example

```
OS10# crypto ca-cert delete Amazon_Root_CA.crt
Successfully removed certificate

OS10# crypto ca-cert delete all
Proceed to delete all installed CA certificates? [confirm yes/no(default)]:yes
```

Supported releases 10.4.3.0 or later

crypto ca-cert install

Installs a certificate from a Certificate Authority that is copied to the switch.

Syntax `crypto ca-cert install` *ca-cet-filepath* `[`*filename*`]`

Parameters

- *ca-cert-filepath* — Enter the local path where the downloaded CA certificate is stored; for example, `home://CAcert.pem` or `usb://CA-cert.pem`.
- *filename* — (Optional) Enter the filename that the CA certificate is stored under in the OS10 trust store directory. Enter the filename in the *filename*`.crt` format.

Default Not configured

Command mode EXEC

Usage information Before using the `crypto ca-cert install` command, copy the certificate to the home directory on the switch using a secure connection, such as HTTPS, SCP, or SFTP. After successful installation, the subject and issuer of the CA certificate are displayed. To delete a trusted certificate, use the `crypto ca-cert delete` command.

Example

```
OS10# crypto ca-cert install home://GeoTrust_Universal_CA.crt
Processing certificate ...
Installed Root CA certificate
CommonName = GeoTrust Universal CA
IssuerName = GeoTrust Universal CA
```

Supported releases 10.4.3.0 or later

crypto cert delete

Deletes an installed host certificate and the private key created with it.

Syntax `crypto cert delete filename [fips]`

Parameters

- *filename* — Enter the file name of the host certificate as displayed in the `show crypto cert` command.
- `fips` — (Optional) Delete a FIPS-compliant certificate-key pair. To verify whether a certificate is non-FIPS or FIPS-compliant., enter the `show crypto cert` command.

Default Not configured

Command mode EXEC

Usage information When you delete the system's certificate, you also delete the private key. Do not delete a host certificate that is used in a security profile. To display the currently installed host certificate and associated key, use the `show crypto cert` command.

NOTE: A FIPS-compliant and non-FIPS certificate may have the same file name. To delete a FIPS-compliant certificate, you must enter the `fips` parameter in the command.

Example

```
OS10# crypto cert delete Dell_host1_CA1.pem
Certificate and keys were successfully deleted. CN = Dell_host1_CA1
```

Supported releases 10.4.3.0 or later

crypto cert generate

Creates a certificate signing request (CSR) or a self-signed certificate.

Syntax `crypto cert generate {request | self-signed} [cert-file cert-path key-file {private | keypath}] [country 2-letter code] [state state] [locality city] [organization organization-name] [orgunit unit-name] [cname common-name] [email email-address] [validity days] [length length] [altname alt-name]`

Parameters

- `request` — Create a certificate signing request to copy to a CA.
- `self-signed` — Create a self-signed certificate.
- `cert-file cert-path` — (Optional) Enter the local path where the self-signed certificate or CSR will be stored. You can enter a full path or a relative path; for example, `flash://certs/s4810-001-request.csr` or `usb://s4810-001.crt`. If you do not enter the `cert-file` option, the system interactively prompts you to fill in the remaining fields of the certificate signing request. Export the CSR to a CA using the `copy` command.

- `key-file {key-path | private}` — Enter the local path where the downloaded or locally generated private key will be stored. If the key was downloaded to a remote server, enter the server path using a secure method, such as HTTPS, SCP, or SFTP. Enter `private` to store the key in a local hidden location.
- `country 2-letter-code` — OPTIONAL) Enter the two letter code that identifies the country.
- `state state` — Enter the name of the state.
- `locality city` — Enter the name of the city.
- `organization organization-name` — Enter the name of the organization.
- `orgunit unit-name` — Enter name of the unit.
- `cname common-name` — Enter the common name assigned to the certificate. Common name is the main identity presented to connecting devices. By default, the switch's host name is the common name. You can configure a different common name for the switch; for example, an IP address. If the *common-name* value does not match the device's presented identity, a signed certificate does not validate.
- `email email-address` — Enter a valid email address used to communicate with the organization.
- `validity days` — Enter the number of days for which the certificate is valid. For a CSR, validity has no effect. For a self-signed certificate, the default is 3650 days or 10 years.
- `length bit-length` — Enter a bit value for the keyword length. For FIPS mode, the range is from 2048 to 4096; for non-FIPS mode, the range is from 1024 to 4096. The default key length for both FIPS and non-FIPS mode is 2048 bits. The minimum key length value for FIPS mode is 2048 bits. The minimum key length value for non-FIPS mode is 1024 bits.
- `altname altname` — Enter an alternate name for the organization; for example, using the IP address such as `altname IP:192.168.1.100`.

Default Not configured

Command mode EXEC

Usage information Generate a CSR when you want a CA to sign a host certificate. Generate a self-signed certificate if you do not set up a CA and implement a certificate trust model in your network.

If you enter the `cert-file` option, you must enter all the following required parameters, including the local path where the certificate and private key are stored.

If you do not specify the `cert-file` option, you are prompted to fill in the other parameter values for the certificate interactively; for example:

```
You are about to be asked to enter information that will be incorporated into
your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank.
For some fields there will be a default value; if you enter '.', the field
will be left blank.
Country Name (2 letter code) [US]:
State or Province Name (full name) [Some-State]:California
Locality Name (eg, city) []:San Francisco
Organization Name (eg, company) []:Starfleet Command
Organizational Unit Name (eg, section) []:NCC-1701A
Common Name (eg, YOUR name) [hostname]:S4148-001
Email Address []:scotty@starfleet.com
```

If the system is in FIPS mode (`crypto fips enable` command), the CSR and private key are generated using approved algorithms from a cryptographic library that has been validated against the FIPS 140-2 standard. You can install the FIPS-compliant certificate-key pair using the `crypto cert install` command with the `fips` option.

Examples

```
OS10# crypto cert generate request cert-file home://cert1.pem key-file home://
cee OS10-VM email admin@dell.com length 1024 altname DNS.dell.com
```

```
Processing certificate ...

Successfully created CSR file /home/admin/cert1.pem and key
```

```
OS10# crypto cert generate self-signed cert-file home://cert2.pem key-file
home:e OS10-VM email admin@dell.com length 1024 altname.dell.com validity 365
Processing certificate ...

Successfully created certificate file /home/admin/cert2.pem and key
```

Supported releases 10.4.3.0 or later

crypto cert install

Installs a host certificate and private key on the switch. A host certificate may be trusted from a CA or self-signed.

Syntax `crypto cert install cert-file` *`cert-path`* `key-file {`*`key-path`* `| private} [password` *`passphrase`*`] [fips]`

Parameters

- `cert-file` *`cert-path`* — Enter the local path to where the downloaded certificate is stored. You can enter a full path or a relative path; for example, `home://s4048-001-cert.pem` or `usb://s4048-001-cert.pem` or `flash://certs/s4810-001-request.crt`.
- `key-file {`*`key-path`* `| private}` — Enter the local path to retrieve the downloaded or locally generated private key. Specify a *key-path* to install the key from a local directory. Enter `private` to install the key from a local hidden location. After the certificate is successfully installed, the private key is deleted from the specified *`key-path`* location and copied to the hidden location.
- `password` *`passphrase`* — (Optional) Enter the password used to decrypt the private key if it was generated using a password.
- `fips` — (Optional) Install the certificate-key pair as FIPS-compliant. Enter `fips` to install a certificate-key pair that a FIPS-aware application, such as RADIUS over TLS, uses. If you do not enter `fips`, the certificate-key pair is stored as a non-FIPS compliant pair.

Default Not configured

Command mode EXEC

Usage information Before using the `crypto cert install` command, copy a CA-signed certificate to the home directory on the switch using a secure connection, such as HTTPS, SCP, or SFTP, and (optionally) the private key. To delete a trusted certificate, use the `crypto cert delete` command.

A successful installation of a trusted certificate requires that:

- The downloaded certificate is correctly formatted.
- The downloaded certificate's public key corresponds to the private key.

You can assign an installed certificate-key pair to a security profile by entering the file name of the certificate without an extension.

It is possible to store a certificate in either FIPS mode or non-FIPS mode on the switch, but not in both modes, using the `crypto cert install` command and the optional `fips` option. You must ensure that certificates installed in FIPS mode are compliant with the FIPS 140-2 standard.

Example

```
OS10# crypto cert install cert-file home://Dell_host1_CA1.pem key-file home://
Dell_host1_CA1.key
Processing certificate ...
Certificate and keys were successfully installed as "Dell_host1_CA1.pem" that
may be used in a security profile. CN = Dell_host1_CA1.
```

Supported releases	10.4.3.0 or later

crypto security-profile

Creates an application-specific security profile.

Syntax	`crypto security-profile` *profile-name*
Parameters	*profile-name* — Enter the name of the security profile, up to 32 characters.
Default	Not configured
Command mode	CONFIGURATION
Usage information	Create a security profile for a specific application on the switch, such as RADIUS over TLS. A security profile associates a certificate and private key pair using the `certificate key` command. The `no` form of the command deletes the security profile.
Example	`OS10# crypto security-profile secure-radius-profile` `OS10(config-sec-profile)#`
Supported releases	10.4.3.0 or later

show crypto ca-certs

Displays all CA certificates installed on the switch.

Syntax	`show crypto ca-certs [`*filename*`]`
Parameters	*filename* — (Optional) Enter the text filename of a CA certificate as shown in the `show crypto ca-certs` output. Enter the filename in the format *filename*`.crt`.
Default	Display all installed CA certificates.
Command mode	EXEC
Usage information	To delete a CA certificate, use the `crypto ca-cert delete` command. Enter the filename as shown in the `show crypto ca-certs` output.
Example	

```
OS10# show crypto ca-certs
 --------------------------------------
|    Locally installed certificates    |
 --------------------------------------
Dell_interCA1.crt
Dell_rootCA1.crt
```

```
OS10# show crypto ca-certs Dell_interCA1.crt
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 4096 (0x1000)
    Signature Algorithm: sha256WithRSAEncryption
        Issuer: C = US, ST = California, L = Santa Clara, O = Dell EMC, OU =
Networking, CN = Dell_rootCA1
        Validity
            Not Before: Jul 25 18:49:22 2018 GMT
            Not After : Jul 22 18:49:22 2028 GMT
        Subject: C = US, ST = California, O = Dell EMC, OU = Networking, CN =
Dell_interCA1
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (4096 bit)
                Modulus:
                    00:b8:46:93:86:27:af:3e:fb:a7:bd:c1:25:76:fd:
                    50:87:02:de:98:2b:95:2e:b0:49:e4:5c:7c:db:83:
```

```
                    b9:e7:3d:e3:61:63:e9:e1:e9:6f:a4:eb:b8:06:bf:
                    57:b7:bb:17:d1:50:ee:7c:ad:d1:09:fe:c3:2c:ea:
                    79:bf:b9:fa:92:52:22:0e:49:62:0b:97:b8:92:c2:
                    59:43:2e:53:e0:c6:d4:ea:d5:ec:35:79:4f:c2:95:
                    82:91:43:ee:3e:3d:ae:e3:a9:ba:37:94:79:27:b3:
                    0d:f9:5a:cc:1b:fd:6d:24:d6:00:ce:1d:3d:4a:fa:
                    95:94:c8:a5:1c:65:cc:f0:08:4a:7f:79:c7:68:4e:
                    c2:3a:b5:b9:21:82:1c:25:45:f4:7e:84:f9:d3:af:
                    28:06:0b:8d:da:72:c1:41:1a:ca:c1:63:de:d6:25:
                    ef:f8:ec:a7:93:88:e0:a0:4f:93:14:81:a6:e8:90:
                    31:7a:b8:53:4c:52:44:e1:5c:6a:aa:94:b6:0d:eb:
                    73:b8:18:21:d5:9c:a4:73:a4:54:16:5b:af:b0:35:
                    0d:36:ff:cb:72:04:63:d1:df:48:59:d3:e9:51:e1:
                    cb:2a:61:20:ee:31:25:51:68:0e:be:98:c3:22:98:
                    29:f9:13:03:c4:2d:bb:4a:d2:cf:7d:00:f9:4c:2e:
                    46:70:e3:ab:e7:3c:91:b0:c9:f7:48:89:ea:e7:df:
                    4f:f4:f5:fc:3a:17:dc:f8:8c:48:e5:aa:03:84:d7:
                    20:7b:55:2e:73:63:85:1c:97:a1:bb:96:95:a1:d3:
                    ae:0c:7a:ae:02:3c:2c:07:b6:9b:c5:97:69:fa:88:
                    bd:ec:8b:88:b3:90:e3:dc:aa:98:15:c6:91:99:a4:
```

Supported releases 10.4.3.0 or later

show crypto cert

Displays information about a specified certificate or all installed certificates.

Syntax `show crypto cert [filename]`

Parameters *filename* — (Optional) Enter the text filename of a certificate as displayed in the `show crypto certs` output. Enter the filename in the format *filename*`.crt`.

Default Display all installed host certificates.

Command mode EXEC

Usage information To delete a certificate, use the `crypto cert delete` *filename* command.

Example

```
OS10# show crypto cert
 --------------------------------------
|    Installed non-FIPS certificates    |
 --------------------------------------
Dell_host1_CA1.pem
 --------------------------------------
|    Installed FIPS certificates    |
 --------------------------------------
```

```
OS10# show crypto cert Dell_host1_CA1.pem
 ------------ Non FIPS certificate -----------------
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 4096 (0x1000)
    Signature Algorithm: sha256WithRSAEncryption
        Issuer: C = US, ST = California, O = Dell EMC, OU = Networking, CN =
Dell_interCA1
        Validity
            Not Before: Jul 25 19:11:19 2018 GMT
            Not After : Jul 22 19:11:19 2028 GMT
        Subject: C = US, ST = California, L = Santa Clara, O = Dell EMC, OU =
Networking, CN = Dell_host1_CA1
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (2048 bit)
                Modulus:
                    00:e7:81:4b:4a:12:8d:ce:88:e6:73:3f:da:19:03:
                    c6:56:01:19:b2:02:61:3f:5b:1e:33:28:a1:ed:e3:
```

```
                    85:bc:56:fb:18:d5:16:2e:a0:e7:3a:f9:34:b4:df:
                    37:97:93:a9:b9:94:b2:9f:69:af:fa:31:77:68:06:
                    89:7b:6d:fc:91:14:4a:c8:7b:23:93:f5:44:5a:0a:
                    3f:ce:9b:af:a6:9b:49:29:fd:fd:cb:34:40:c4:02:
                    30:95:37:28:50:d8:81:fb:1f:83:88:d9:1f:a3:0e:
                    49:a1:b3:df:90:15:d4:98:2b:b2:38:98:6e:04:aa:
                    bd:92:1b:98:48:4d:08:49:69:41:4e:6a:ee:63:d8:
                    2a:9f:e6:15:e2:1d:c3:89:f5:f0:d0:fb:c1:9c:46:
                    92:a9:37:b9:2f:a0:73:cf:e7:d1:88:96:b8:4a:84:
                    91:83:8c:f0:9a:e0:8c:6e:7a:fa:6e:7e:99:3a:c3:
                    2c:04:f9:06:8e:05:21:5f:aa:6e:9f:b7:10:37:29:
                    0c:03:14:a0:9d:73:1f:95:41:39:9b:96:30:9d:0a:
                    cb:d0:65:c3:59:23:01:f7:f5:3a:33:b9:e9:95:11:
                    0c:51:f4:e9:1e:a5:9d:f7:95:84:9c:25:74:0c:21:
                    4f:8b:07:29:2f:e3:47:14:50:8b:03:c1:fb:83:85:
                    dc:bb
                Exponent: 65537 (0x10001)
        X509v3 extensions:
            X509v3 Basic Constraints:
                CA:FALSE
            Netscape Cert Type:
                SSL Client, S/MIME
            Netscape Comment:
                OpenSSL Generated Client Certificate
            X509v3 Subject Key Identifier:
                4A:20:AA:E1:69:BF:BE:C5:66:2E:22:71:70:B4:7E:32:6F:E0:05:28
            X509v3 Authority Key Identifier:
                keyid:A3:39:CB:C7:76:86:3B:05:44:34:C2:6F:90:73:1F:5F:64:55:5C:
76
            X509v3 Key Usage: critical
```

Supported releases 10.4.3.0 or later

Example: Configure RADIUS over TLS with X.509v3 certificates

This example shows how to install a trusted X.509v3 CA and a host certificate-key pair that supports RADIUS over TLS authentication.

1. Install a trusted CA certificate.

```
OS10# copy tftp://CAadmin:secret@172.11.222.1/GeoTrust_Universal_CA.crt home://
GeoTrust_Universal_CA.crt
OS10# crypto ca-cert install home://GeoTrust_Universal_CA.crt
Processing certificate ...
Installed Root CA certificate
CommonName = GeoTrust Universal CA
IssuerName = GeoTrust Universal CA
```

2. Generate a CSR, copy the CSR to a CA server, download the signed certificate, and install the host certificate.

```
OS10# crypto cert generate request cert-file home://s4048-001-csr.pem
key-file home://tsr6-key.pem cname "Top of Rack 6" altname "IP:10.0.0.6 DNS:tor6.dell.com"
email admin@dell.com organization "Dell EMC" orgunit Networking locality "santa Clara"
state California country US length 1024
Processing certificate ...
Successfully created CSR file /home/admin/tor6-csr.pem and key

OS10# copy home://tor6-csr.pem scp://CAadmin:secret@172.11.222.1/s4048-001-csr.pem

OS10# copy scp://CAadmin:secret@172.11.222.1/s4048-001.crt usb://s4048-001-crt.pem

OS10# crypto cert install crt-file usb://s4048-001-crt.pem key-file usb://s4048-001-crt.key
This will replace the already installed host certificate.
Do you want to proceed ? [yes/no(default)]:yes
Processing certificate ...
Host certificate installed successfully.
```

3. Configure an X.509v3 security profile.

```
OS10# show crypto cert
-------------------------------------
|    Installed non-FIPS certificates    |
-------------------------------------
s4048-001-csr.pem
-------------------------------------
|    Installed FIPS certificates    |
-------------------------------------

OS10# config terminal
OS10(config)# crypto security-profile radius-admin
OS10(config-sec-profile)# certificate s4048-001-csr
OS10(config-sec-profile)# exit
```

4. Configure the RADIUS over TLS server.

```
OS10# radius-server host 10.0.0.1 tls security-profile radius-admin key radsec
```

5. Configure RADIUS-based user authentication.

```
OS10# aaa authentication login default group radius local
```

Simple Network Management Protocol

Network management stations use simple network management protocol (SNMP) to retrieve and modify software configurations for managed objects on an agent in network devices. A *managed object* is a datum of management information.

The SNMP agent in a managed device maintains the data for managed objects in management information bases (MIBs). Managed objects are identified by their object identifiers (OIDs). A remote SNMP agent performs an SNMP walk on the OIDs stored in MIBs on the local switch to view and retrieve information.

OS10 supports standard and private SNMP MIBs, including all `get` requests. MIBs are hierarchically structured and use object identifiers to access managed objects. For a list of MIBs supported in the OS10 version running on a switch, refer to the *OS10 Release Notes* for the release.

NOTE: OS10 supports SNMP set operations only on SysName in System MIBs.

OS10 supports different security models and levels in SNMP communication between SNMP managers and agents. Each security model refers to an SNMP version used in SNMP messages. SNMP versions provide different levels of security, such as user authentication and message encryption.

MIBs

Release 10.4.3.0 supports the following standard and Dell EMC enterprise MIBs.
The MIBs are stored in the following folder: **/opt/dell/os10/snmp/mibs/**.

Standards MIBs:

Table 32. Standards MIBs

Module	Standard
BRIDGE-MIB	IEEE 802.1D
ENTITY-MIB	RFC 6933
EtherLike-MIB	RFC 3635

Module	Standard
HOST-RESOURCES-MIB	RFC 2790
IEEE8021-PFC-MIB	IEEE 802.1Qbb
IEEE8023-LAG-MIB	IEEE 802.3ad
IF-MIB	RFC 2863
IP-FORWARD-MIB	RFC 4292
IP-MIB	RFC 4293
LLDP-EXT-DOT1-MIB	IEEE 802.1AB
LLDP-EXT-DOT3-MIB	IEEE 802.1AB
LLDP-MIB	IEEE 802.1AB
OSPF-MIB	RFC 4750
OSPFV3-MIB	RFC 5643
Q-BRIDGE-MIB	IEEE 802.1Q
RFC1213-MIB	RFC 1213
SFLOW-MIB	RFC 3176
SNMP-FRAMEWORK-MIB	RFC 3411
SNMP-MPD-MIB	RFC 3412
SNMP-TARGET-MIB	RFC 3413
SNMP-USER-BASED-SM-MIB	RFC 3414
SNMP-VIEW-BASED-ACM-MIB	RFC 3415
SNMPv2-MIB	RFC 3418
TCP-MIB	RFC 4022
UDP-MIB	RFC 4113

Dell EMC Enterprise MIBs:

Table 33. Dell EMC Enterprise MIBs

Module	Description
DELLEMC-OS10-BGP4V2-MIB	OS10 BGPv2 implementations
DELLEMC-OS10-CHASSIS-MIB	OS10 chassis implementations
DELLEMC-OS10-PRODUCTS-MIB	OS10 platform product definitions
DELLEMC-OS10-SMI-MIB	OS10 SMI implementations
DELLEMC-OS10-TC-MIB	OS10 networking equipment textual convention

SNMP security models and levels

OS10 supports SNMP security models v1, v2c, and v3. The supported security levels are no authentication, authentication, and privacy.

You specify the SNMP security model and level when you configure SNMP groups and users. Each security model corresponds to an SNMP version that provides different security levels:

- SNMPv1 provides no user authentication or privacy protection (encryption). SNMP messages are sent in plain text.
- SNMPv2c provides no user authentication or encryption. SNMP messages are sent in plain text.
- SNMPv3 provides user-configured security levels for user authentication and encryption of SNMP messages:
  - No user password or message encryption
  - User authentication only
  - User authentication and message encryption

SNMPv3

SNMP version 3 (SNMPv3) provides an enhanced security model for user authentication and encryption of SNMP messages. User authentication requires that SNMP packets come from an authorized source. Message encryption ensures that packet contents cannot be viewed by an unauthorized source.

To configure SNMPv3-specific security settings — user authentication and message encryption — use the `snmp-server user` command. You can generate localized keys with enhanced security for authentication and privacy (encryption) passwords.

SNMP engine ID

An engine ID identifies the SNMP entity, local agent, on the switch. The engine ID is an octet colon-separated number; for example, `00:00:17:8B:02:00:00:01`.

When you configure an SNMPv3 user, you can specify that a localized authentication and/or privacy key be generated. The localized password keys are generated using the engine ID of the switch. A localized key is more complex and provides greater privacy protection.

The engine ID used to generate the password keys is unique to the switch. For this reason, you cannot copy and use localized SNMP security passwords on another switch.

SNMP groups and users

A member of an SNMP group that accesses the local SNMP agent is referred to as an *SNMP user*. An SNMP user on a remote device is identified by an IP address and UDP port from which the user accesses the local agent.

In OS10, users are assigned SNMP access privielges according to the group they belong to. You configure each group for access to SNMP MIB tree views.

SNMP views

In OS10, you configure views for each security model and level in an SNMP user group. Each type of view specifies the object ID (OID) in the MIB tree hierarchy at which the view starts. You can also specify whether the rest of the MIB tree structure is included or excluded from the view.

- A *read* view provides read-only access to the specified OID tree.
- A *write* view provides read-write access to the specified OID tree.
- A *notify* view allows SNMP notifications (traps and informs) from the specified OID tree to be sent to other members of the group.

Configure SNMP

To set up communication with SNMP agents in your network:

- Configure the read-only, read-write, and notify access for SNMP groups.
- Configure groups with SNMP views for specified SNMP versions (security models).
- Assign users to groups and configure SNMPv3-specific authentication and encryption settings, and optionally, localized security keys and ACL-based access.

Configuring SNMP consists of these tasks in any order:

- Configure SNMP engine ID
- Configure SNMP views
- Configure SNMP groups
- Configure SNMP users

Configure SNMP engine ID

The engine ID identifies the SNMP local agent on a switch. The engine ID is an octet colon-separated number; for example, `80:00:02:b8:04:61:62:63`.

The local engine ID is used to create a localized authentication and/or privacy key for greater security in SNMPv3 messages. You generate a localized authentication and/or privacy key when you configure an SNMPv3 user.

Configure a remote device and its engine ID to allow a remote user to query the local SNMP agent. The remote engine ID is included in the query and used to generate the authentication and privacy password keys to access the local agent. If you do not configure the remote engine ID, remote users' attempts to access the local agent fail.

NOTE: Be sure to create a remote engine ID with the `snmp-server engineID` command before you configure a remote user with the `snmp-server user` command. If you change the configured engine ID for a remote device, you must reconfigure the authentication and privacy passwords for all remote users associated with the remote engine ID.

```
snmp-server engineID [local engineID] [remote ip-address {[udp-port port-number] remote-
engineID}]
```

To display the localized authentication and privacy keys in an SNMPv3 user configuration, enter the `show snmp engineID local` command.

Generate SNMPv3 localized keys

```
OS10(config)# snmp-server engineID local 80:00:02:b8:04:61:62:63

OS10(config)# snmp-server engineID remote 1.1.1.2 udp-port 432 0xabeecc
```

Display localized keys

```
OS10# show snmp-server engineID local
Local default SNMP engineID: 80:00:02:b8:04:61:62:63
```

Configure SNMP views

Configure a read-only, read-write, or notify view of the MIB tree structure in the SNMP agent on the switch.

The `oid-tree` value specifies the OID in the MIB tree hierarchy at which a view starts. Enter `included` or `excluded` to include or exclude the rest of the sub-tree MIB contents in the view. If necessary, re-enter the command to exclude tree entries in the included content.

```
snmp-server view view-name oid-tree [included | excluded]
```

Configure read-only view

```
OS10(config)# snmp-server view readonly 1.3.6.1.2.1.31.1.1.1.6 included
```

Configure read-write view

```
OS10(config)# snmp-server view rwView 1.3.6.1.2.1.31.1.1.1.6 included
OS10(config)# snmp-server view rwView 1.3.6.1.2.1.31.0.0.0.0 excluded
```

Display SNMP views

```
OS10# show snmp-server view
view name               : readview
OID                     : 1.3.6.5
excluded                : True
```

Configure SNMP groups

Configure an SNMP group with the views allowed for the members of the group. Specify the read-only, read-write, and/or notification access to the SNMP agent.

The security model corresponds to the SNMP version that users use to send and receive SNMP messages. The security level configures SNMPv3 user authentication and privacy settings:

- `auth` — Authenticate users in SNMP messages.
- `noauth` — Do not authenticate users or encrypt SNMP messages; send messages in plain text.
- `priv` — Authenticate users and encrypt/decrypt SNMP messages.

Enter an ACL to limit user access so that only messages from and to ACL-allowed users are received and sent from the SNMP agent on the switch.

```
snmp-server group group-name {v1 | v2c | v3 security-level} [access acl-name]
[read view-name] [write view-name] [notify view-name]
```

To configure a view of the MIB tree on the SNMP agent, use the `snmp-server view` command.

To configure an SNMPv3 user's authentication and privacy settings, use the `snmp-server user` command.

To display the configured SNMP groups, use the `show snmp group` command.

Configure SNMPv1 or v2c group

```
OS10(config)# snmp-server group v2group 2c read readview notify GetsSets
```

Configure SNMPv3 group

```
OS10(config)# snmp-server group v3group 3 priv read readview write writeview notify alltraps
```

Display SNMP groups

```
OS10# show snmp-server group
groupname               : v2group
version                 : 2c
notifyview              : GetsSets
readview                : readview

groupname               : v3group
```

```
version                 : 3
security level          : priv
notifyview              : alltraps
readview                : readview
writeview               : writeview
```

Configure SNMP users

Configure user access to the SNMP agent on the switch using group membership. Assign each user to a group and configure SNMPv3-specific authentication and encryption settings, and optionally, localized security keys and ACL-based access. Re-enter the command multiple times to configure SNMP security settings for all users.

```
snmp-server user user-name group-name security-model [[noauth | auth {md5 | sha} auth-password]
[priv {des | aes}]] [localized] [access acl-name]  [remote ip-address udp-port port-number]]
```

The group to which a user is assigned determines the user's access privilege. To configure a group's access privilege — read, write, and notify — to the switch, use the `snmp-server group` command. The security model for SNMPv3 provides the strongest security with user authentication and packet encryption.

No default values exist for SNMPv3 authentication and privacy algorithms and passwords. If you forget a password, you cannot recover it — you must reconfigure the user. You can specify either a plain-text password or an encrypted cypher-text password. In either case, the password stores in the configuration in encrypted form and displays as encrypted in `show running-config snmp` output.

A localized authentication or privacy key is more complex and provides greater privacy protection. Localized keys are generated using the engine ID of the switch. For this reason, you cannot use the localized SNMP security passwords in the configuration file on another switch. For more information, see Configure SNMP engine ID. To display the localized authentication and privacy keys in an SNMPv3 user configuration, use the `show running-configuration snmp` command.

To limit user access to the SNMP agent on the switch, enter an `access acl-name` value. In IPv6 ACLs, SNMP supports only IPv6 and UDP types. TCP, ICMP, and port rules are not supported.

To display the configured SNMP users, use the `show snmp user` command.

Configure SNMPv1 or v2c users

```
OS10(config)# snmp-server user admin1 netadmingroup 2c acl acl_AdminOnly
```

Configure SNMPv3 users

```
OS10(config)# snmp-server user privuser v3group 3 encrypted auth
md59fc53d9d908118b2804fe80e3ba8763d priv des56 d0452401a8c3ce42804fe80e3ba8763d

OS10(config)# snmp-server user n3user ngroup remote 172.31.1.3 udp-port 5009 3auth md5
authpasswd
```

Display SNMP users

```
OS10# show snmp-server user
User name                : privuser
Group                    : v3group
Version                  : 3
Authentication Protocol  : MD5
Privacy Protocol         : AES
```

SNMP commands

SNMP traps: Enable SNMP notifications to send to network management host devices.

show snmp community

Displays the SNMP communities configured on the switch.

Syntax `show snmp community`

Parameters None

Defaults None

Command Mode EXEC

Usage Information To configure an SNMP community, use the `snmp-server community` command.

Example

```
OS10# show snmp community
Community      : public
Access         : read-only

Community      : dellOS10
Access         : read-write
ACL            : dellacl
```

Supported Releases 10.4.2.0 or later

show snmp engineID

Displays the SNMP engine ID on the switch or on remote devices that access the SNMP agent on the switch.

Syntax `show snmp engineID {local | remote}`

Parameters

- `local` — Display the local engine ID.
- `remote` — Display the SNMP engine ID of remote devices configured on the switch.

Defaults None

Command Mode EXEC

Usage Information To configure the local engine ID or the engine ID for a remote device, use the `snmp-server engineID` command.

Example

```
OS10# show snmp engineID remote
Remote Engine ID     IP-addr     Port
0x0712               1.1.1.1     23

OS10# show snmp engineID local
Local default SNMP engineID: 0x80001f880390b11cf4abe7
```

Supported Releases 10.4.2.0 or later

show snmp group

Displays the SNMP groups configured on the switch, including SNMP views and security models.

Syntax `show snmp group`

Parameters None

Defaults None

Command Mode EXEC

Usage Information To configure an SNMP group, use the `snmp-server group` command.

Example

```
OS10# show snmp group
groupname                : v2group
version                  : 2c
notifyview               : GetsSets
readview                 : readview

groupname                : v3group
version                  : 3
security level           : priv
notifyview               : alltraps
readview                 : readview
writeview                : writeview
```

Supported Releases 10.4.2.0 or later

show snmp user

Displays the users configured to access the SNMP agent on the switch, including the SNMP group and security model.

Syntax `show snmp user`

Parameters None

Defaults None

Command Mode EXEC

Usage Information To configure an SNMP user, use the `snmp-server user` command.

Example

```
OS10# show snmp user
User name                : privuser
Group                    : v3group
Version                  : 3
Authentication Protocol  : MD5
Privacy Protocol         : AES
```

Supported Releases 10.4.2.0 or later

show snmp view

Displays the SNMP views configured on the switch, including the SNMP object ID at which the view starts.

Syntax `show snmp view`

Parameters None

Defaults None

Command Mode EXEC

Usage Information Use the `show snmp view` command to verify the OID starting point for SNMP views in MIB trees. To configure an SNMP view, use the `snmp-server view` command.

Example

```
OS10# show snmp view
view name                : readview
```

```
OID                     : 1.3.6.5
excluded                : True
```

Supported Releases 10.4.2.0 or later

snmp-server community

Configures an SNMP user community.

Syntax `snmp-server community name {ro | rw} [acl acl-name]`

Parameters

- `community name` — Set the community name string to act as a password for SNMPv1 and SNMPv2c access. A maximum of 20 alphanumeric characters.
- `ro` — Set read-only access for the SNMP community.
- `rw` — Set read-write access for the SNMP community.
- `acl acl-name` — Enter an existing IPv4 ACL name to limit SNMP access in the SNMP community.

Defaults An SNMP community has read-only access.

Command Mode CONFIGURATION

Usage Information The SNMPv1 and SNMPv2c security models use a community-based form of security. Use the `snmp-server community` command to configure read-only or read-write access for an SNMP community name. The configured community text string is used for SNMPv1 and SNMPv2c user authentication.

To display the SNMP communities on the switch, use the `show snmp-server community` command. The `no` version of the command removes the configured community text string.

Example

```
OS10(config)# snmp-server community admin rw
OS10(config)# snmp-server community public ro acl snmp-read-only-acl
```

Supported Releases 10.2.0E or later

snmp-server contact

Configures contact information for troubleshooting the local SNMP switch.

Syntax `snmp-server contact text`

Parameters `text` — Enter an alphanumeric text string. A maximum of 55 characters.

Default The SNMP server contact is `support`.

Command Mode CONFIGURATION

Usage Information The `no` version of this command resets the SNMP server contact to the default value.

Example

```
OS10(config)# snmp-server contact administrator
```

Supported Releases 10.2.0E or later

snmp-server enable traps

Enables SNMP traps on a switch.

Syntax `snmp-server enable traps [notification-type] [notification-option]`

Parameters

- *notification-type notification-option* — Enter an SNMP notification type, and optionally, a notification option for the type.

Table 34. Notification types and options

Notification type	Notification option
`entity` — Enable entity change traps.	None
`envmon` — Enable SNMP environmental monitor traps.	– `fan` — Enable fan traps. – `power-supply` — Enable power-supply traps. – `temperature` — Enable temperature traps.
`lldp` — Enable LLDP state change traps.	– `rem-tables-change` — Enable the lldpRemTablesChange trap.
`snmp` — Enable SNMP traps.	– `authentication` — Enable authentication traps. – `coldstart` — Enable coldstart traps when you power on the switch and the SNMP agent initializes. – `linkdown` — Enable link-down traps. – `linkup` — Enable link-up traps. – `warmstart` — Enable warmstart traps when the switch reloads and the SNMP agent reinitializes.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information If you do not enter a `notification-type` or `notification-option` parameter with the command, all traps are enabled. If you enter only a `notification-type`, all `notification-option` traps associated with the type are enabled.

Re-enter the command multiple times with different notification types and options to enable specific SNMP trap types.

The `no` version of the `snmp-server enable traps` command disables SNMP traps on the switch.

Example

```
OS10(config)# snmp-server enable traps envmon fan
OS10(config)# snmp-server enable traps envmon power-supply
OS10(config)# snmp-server enable traps snmp

OS10(config)# no snmp-server enable traps snmp
```

Supported Releases 10.4.1.0 or later

snmp-server engineID

Configures the local and remote SNMP engine IDs.

Syntax

```
snmp-server engineID [local engineID] [remote ip-address {[udp-port port-number] remote-engineID}]
```

Parameters

- `local engineID` — Enter the engine ID that identifies the local SNMP agent on the switch as an octet colon-separated number. A maximum of 27 characters.
- `remote ip-address` — Enter the IPv4 or IPv6 address of a remote SNMP device that accesses the local SNMP agent.
- `udp-port port-number` — Enter the UDP port number on the remote device, from 0 to 65535.
- `remote-engineID` — Enter the engine ID that identifies the SNMP agent on a remote device, `0x` followed by a hexadecimal string).

Defaults

The local engine ID is generated using the MAC address of the management Ethernet interface.

Command Mode

CONFIGURATION

Usage Information

The local engine ID is used to generate the localized keys for the authentication and privilege passwords. These passwords authenticate SNMP users and encrypt SNMP messages. If you reconfigure the local Engine ID, the localized keys also change. The existing values are no longer valid and a warning message is displayed. As a result, you must reconfigure SNMP users with new localized password keys.

In addition, if you change the configured engine ID for a remote device, you must reconfigure the authentication and privacy passwords for the remote user.

To display the current local engine ID, use the `show snmp engineID local` command. The `no` version of this command resets the default engine ID values.

Example

```
OS10(config)# snmp-server engineID local 80:00:02:b8:04:61:62:63

OS10(config)# snmp-server engineID local 80:00:02:b8:04:61:62:63
% Warning: Localized passwords need to be regenerated for local user.

OS10(config)# snmp-server engineID remote 1.1.1.1 0xaaffcc

OS10(config)# snmp-server engineID remote 1.1.1.2 udp-port 432 0xabeecc
```

Supported Releases

10.4.2.0 or later

snmp-server group

Configures the views allowed for the users in an SNMP group.

Syntax

```
snmp-server group group-name {v1 | v2c | v3 security-level} [access acl-name] [read view-name] [write view-name] [notify view-name]
```

Parameters

- `group-name` — Enter the name of the group. A maximum of 32 alphanumeric characters.
- `v1` — SNMPv1 provides no user authentication or privacy protection. SNMP messages are sent in plain text.
- `v2c` — SNMPv2c provides no user authentication or privacy protection. SNMP messages are sent in plain text.

- `v3 security-level` — SNMPv3 provides optional user authentication and encryption for SNMP messages, configured with the `snmp-server user` command.
- `security-level` — (SNMPv3 only) Configure the security level for SNMPv3 users:
  - `auth` — Authenticate users in SNMP messages.
  - `noauth` — Do not authenticate users or encrypt SNMP messages; send messages in plain text.
  - `priv` — Authenticate users and encrypt/decrypt SNMP messages.
- `access acl-name` — (Optional) Enter the name of an IPv4 or IPv6 access list to filter SNMP requests received on the switch. A maximum of 16 characters.
- `read view-name` — (Optional) Enter the name of a read-only view. A maximum of 32 characters maximum.
- `write view-name` — (Optional) Enter the name of a read-write view. A maximum of 32 characters maximum.
- `notify view-name` — (Optional) Enter the name of a notification view. A maximum of 32 characters maximum.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information Use the `snmp-server group` command to set up the access privileges for a group of SNMP users. Configure the security level for receiving SNMP messages. Specify read-only, read-write, and/or notification access to the SNMP agent. To configure an SNMPv3 user's authentication and privacy settings, use the `snmp-server user` command.

Enter an `access acl-name` value to limit access to the SNMP agent on the switch to only ACL-allowed users.

A read-view provides read-only access to the SNMP agent. A read-write view allows read-write access. A notify-view allows SNMP notifications to be sent to group members.

The `no snmp-server group group-name` command deletes an SNMP group.

Example

```
OS10(config)# snmp-server group os10admin p3 priv read readonlyview
```

Supported Releases 10.4.2.0 or later

snmp-server host

Configures a host to receive SNMP notifications.

Syntax

```
snmp-server host {ipv4-address | ipv6-address} {informs version version-number
| traps version version-number | version version-number} [snmpv3-security-
level] [community-name] [udp-port port-number] [entity | envmon | lldp | snmp]
```

Parameters

- `ipv4-address | ipv6-address` — Enter the IPv4 or IPv6 address of the SNMP host.
- `informs` — Send inform messages to the SNMP host
- `traps` — Send trap messages to the SNMP host
- `version version-number` — Enter the SNMP security model used to send traps or informs to the SNMP host — `1`, `2c`, or `3`. All security models support traps; only `2c` and `3` support informs. To send only SNMP notifications, enter only a *version-number*; do not enter `informs` or `traps`. For SNMPv3 traps and informs, enter the security level:
  - `noauth` — (SNMPv3 only) Send SNMPv3 traps without user authentication and privacy encryption.
  - `auth` — (SNMPv3 only) Include a user authentication key for SNMPv3 messages sent to the host:
    - `md5` — Generate an authentication key using the Message Digest Algorithm (MD5) algorithm.
    - `sha` — Generate an authentication key using the Secure Hash Algorithm (SHA) algorithm.

- *auth-password* — Enter a text string used to generate the authentication key that identifies the user. A maximum of 32 alphanumeric characters. For an encrypted password, enter the encrypted string instead of plain text.

– priv — (SNMPv3 only) Configure encryption for SNMPv3 messages sent to the host:

- aes — Encrypt messages using an AES 128-bit algorithm.
- des — Encrypt messages using a DES 56-bit algorithm.
- *priv-password* — Enter a text string used to generate the privacy key used in encrypted messages. A maximum of 32 alphanumeric characters. For an encrypted password, you can enter the encrypted string instead of plain text.

- *community-name* — (Optional) Enter an SNMPv1 or SNMPv2c community string name or an SNMPv3 user name.
- udp-port *port-number* — (Optional) Enter the UDP port number on the SNMP host, from 0 to 65535.
- entity | envmon | lldp | snmp — Enter one or more types of traps and notifications to send to the SNMP host — entity change, environment monitor, or LLDP state change traps, or SNMP-type notifications.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The local SNMP agent sends SNMP notifications, traps, and informs to SNMP managers configured as host receivers. You can configure multiple host receivers.

An SNMP host does not acknowledge the trap messages and notifications received from the SNMP agent. SNMP hosts send an acknowledgement when receiving informs.

The no version of this command disables the local agent from sending SNMP traps, informs, or notifications to a host receiver.

Example — Send SNMP traps to host

```
OS10(config)# snmp-server host 1.1.1.1 traps version 3 priv user01 udp-port 32 entity lldp
```

Example — Send SNMP informs to host

```
OS10(config)# snmp-server host 1.1.1.1 informs version 2c public envmon snmp
```

Example — Send SNMP notifications to host

```
OS10(config)# snmp-server host 1.1.1.1 version 3 noauth u1 snmp lldp
```

Supported Releases 10.2.0E or later

snmp-server location

Configures the location of the SNMP server.

Syntax snmp-server location *text*

Parameters *text* — Enter an alphanumeric string. A maximum of 55 characters.

Default None

Command Mode CONFIGURATION

Usage Information The no version of this command removes the SNMP location.

Example

```
OS10(config)# snmp-server location datacenter10
```

Supported Releases 10.2.0E or later

snmp-server user

Authorizes a user to access the SNMP agent and receive SNMP messages.

Syntax

```
snmp-server user user-name group-name security-model [[noauth | auth {md5 | sha} auth-password] [priv {des | aes} priv-password]] [localized] [access acl-name] [remote ip-address udp-port port-number]]
```

Parameters

- *user-name* — Enter the name of the user. A maximum of 32 alphanumeric characters.
- *group-name* — Enter the name of the group to which the user belongs. A maximum of 32 alphanumeric characters.
- *security-model* — Enter an SNMP version that sets the security level for SNMP messages:
  - 1 — SNMPv1 provides no user authentication or privacy protection. SNMP messages are sent in plain text.
  - 2c — SNMPv2c provides no user authentication or privacy protection. SNMP messages are sent in plain text.
  - 3 — SNMPv3 provides optional user authentication and encryption for SNMP messages.
- noauth — (SNMPv3 only) Configure SNMPv3 messages to send without user authentication and privacy encryption.
- auth — (SNMPv3 only) Include a user authentication key for SNMPv3 messages sent to the user:
  - md5 — Generate an authentication key using the MD5 algorithm.
  - sha — Generate an authentication key using the SHA algorithm.
  - *auth-password* — Enter a text string used to generate the authentication key that identifies the user (32 alphanumeric characters maximum). For an encrypted password, you can enter the encrypted string instead of plain text.
- priv — (SNMPv3 only) Configure encryption for SNMPv3 messages sent to the user:
  - aes — Encrypt messages using AES 128-bit algorithm.
  - des — Encrypt messages using DES 56-bit algorithm.
  - *priv-password* — Enter a text string used to generate the privacy key used in encrypted messages. A maximum of 32 alphanumeric characters. For an encrypted password, enter the encrypted string instead of plain text.
- localized — (SNMPv3 only) Generate an SNMPv3 authentication and/or privacy key in localized key format.
- access *acl-name* — (Optional) Enter the name of an IPv4 or IPv6 access list to filter SNMP requests on the switch. A maximum of 16 characters.
- remote *ip-address/prefix-length* udp-port *port-number* — (Optional) Enter the IPv4 or IPv6 address of the user's remote device and the UDP port number used to connect to the SNMP agent on the switch, from 0 to 65535; default 162.

Defaults

Not configured

Command Mode

CONFIGURATION

Usage Information

Use the snmp-server user command to set up the desired security level for SNMP access. For SNMPv3 users, configure user authorization and message encryption. Re-enter this command multiple times to configure SNMP security settings for all users.

The group to which a user is assigned determines the user's SNMP access. To configure a group's SNMP access to the switch — read, write, and notify, use the snmp-server user command.

No default values exist for SNMPv3 authentication and privacy algorithms and passwords. If you forget a password, you cannot recover it — you must reconfigure the user. You can specify either a plain-text password or

an encrypted cypher-text password. In either case, the password stores in the configuration in an encrypted form and displays as encrypted in the `show running-config snmp` output.

A localized authentication or privacy key is more complex and provides greater privacy protection. To display the localized authentication and privacy keys in an SNMPv3 user configuration, use the `show running-configuration snmp` command.

To limit user access to the SNMP agent on the switch, enter an `access` *`acl-name`* value. In IPv6 ACLs, SNMP supports only IPv6 and UDP types. TCP, ICMP, and port rules are not supported.

The `no` version of this command removes a user from the SNMP group.

Example (Encrypted passwords)

```
OS10(config)# snmp-server user privuser v3group v3 auth md5
9fc53d9d908118b2804fe80e3ba8763d priv des d0452401a8c3ce42804fe80e3ba8763d
```

Example (Plain-text passwords)

```
OS10(config)# snmp-server user authuser v3group v3 auth md5 authpasswd
```

Example (Remote user)

```
OS10(config)# snmp-server user n3user ngroup remote 172.31.1.3 udp-port 5009 3
auth md5 authpasswd
```

Supported Releases 10.4.2.0 or later

snmp-server view

Configures an SNMPv3 view.

Syntax

```
snmp-server view view-name oid-tree [included | excluded]
```

Parameters

- *`view-name`* — Enter the name of a read-only, read-write, or notify view. A maximum of 32 characters.
- *`oid-tree`* — Enter the SNMP object ID at which the view starts in 12-octet dotted-decimal format.
- `included` — (Optional) Include the MIB family in the view.
- `excluded` — (Optional) Exclude the MIB family from the view.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The `oid-tree` value specifies the OID in the MIB tree hierarchy at which a view starts. Enter `included` or `excluded` to include or exclude the remaining part of the MIB sub-tree contents in the view.

The `no` version of this command removes an SNMPv3 view.

Example

```
OS10(config)# snmp-server view readview 1.3.6.5 excluded
```

Supported Releases 10.4.2.0 or later

snmp-server vrf

Configures an SNMP agent to receive SNMP traps for the management VRF instance.

Syntax

```
snmp-server vrf management
```

Parameters	None
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command disables the SNMP agent from receiving the SNMP traps.
Example	`OS10(config)# snmp-server vrf management`
Supported Releases	10.4.1.0 or later

OS10 image upgrade

The `image download` command simply downloads the software image — it does not install the software on your device. The `image install` command installs the downloaded image to the standby partition.

NOTE: If the active partition contains any modified text files or custom packages installed, they would not be available in the standby partition. Backup the modified files and re-install the packages after downloading the image.

1. (Optional) Backup the current running configuration to the startup configuration in EXEC mode.
   ```
   copy running-configuration startup-configuration
   ```
2. Backup the startup configuration in EXEC mode.
   ```
   copy config://startup.xml config://<backup file name>
   ```
3. Download the new software image from dell.com/support, extract the `bin` files from the `tar` file, and save the file in EXEC mode.
   ```
   image download file-url
   ```
4. (Optional) View the current software download status in EXEC mode.
   ```
   show image status
   ```
5. Install the software image in EXEC mode.
   ```
   image install image-url
   ```
6. (Optional) View the status of the current software install in EXEC mode. For the S5148F-ON platform, open a new SSH or Telnet session to check the status of the current software.
   ```
   show image status
   ```
7. Change the next boot partition to the standby partition in EXEC mode. Use the `active` parameter to set the next boot partition from standby to active.
   ```
   boot system standby
   ```
8. (Optional) Check whether the next boot partition has changed to standby in EXEC mode.
   ```
   show boot detail
   ```
9. Reload the new software image in EXEC mode.
   ```
   reload
   ```

Image download

```
OS10# image download ftp://userid:passwd@hostip:/filepath
```

Image install

```
OS10# image install image://filename.bin
```

Show version

```
TR2# show version
Dell EMC Networking OS10 Enterprise
Copyright (c) 1999-2019 by Dell Inc. All Rights Reserved.
OS Version: 10.4.3.0
Build Version: 10.4.3.85
Build Time: 2019-02-18T17:06:10-0800
System Type: S4048-ON
```

```
Architecture: x86_64
Up Time: 2 days 05:58:01
```

Boot system partition

Set the boot partition to active or standby for subsequent boot cycles. Boot OS10 from standby to load the image on the standby partition, or boot from active to load the currently running image.

1 Display current boot information in EXEC mode.

```
show boot detail
```

2 Configure the boot system in EXEC mode.

```
boot system [active | standby]
```

- `active` — Resets the running partition as the subsequent boot partition.
- `standby` — Sets the standby partition as the subsequent boot partition.

View boot detail

```
OS10# show boot detail
Current system image information detail:
=========================================
Type:                     Node-id 1
Boot Type:                Flash Boot
Active Partition:         A
Active SW Version:        10.4.3E
Active SW Build Version:  10.4.3E.85
Active Kernel Version:    Linux 4.9.110
Active Build Date/Time:   2019-02-18T09:06:10Z
Standby Partition:        B
Standby SW Version:       10.4.3E
Standby SW Build Version: 10.4.3E.80
Standby Build Date/Time:  2019-02-17T15:36:08Z
Next-Boot:                active[A]
```

View boot summary

```
OS10# show boot
Current system image information:
==================================
Type      Boot Type      Active                   Standby                  Next-Boot
-----------------------------------------------------------------------------------
Node-id 1 Flash Boot     [A] 10.4.3E             [B] 10.4.3E              [A] active
```

Upgrade commands

boot system

Sets the boot partition to use during the next reboot.

Syntax	`boot system {active | standby}`
Parameters	• `active` — Reset the running partition as the next boot partition. • `standby` — Set the standby partition as the next boot partition.
Default	Active

Command Mode	EXEC
Usage Information	Use this command to configure the location of the OS10 image used to reload the software at boot time. Use the `show boot` command to view the configured next boot image. This command applies immediately and does not require the `commit` command.
Example	`OS10# boot system standby`
Supported Releases	10.2.0E or later

image cancel

Cancels an active image download.

Syntax	`image cancel`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	This command attempts to cancel an active file download in progress.
Example	`OS10# image cancel`
Supported Releases	10.2.0E or later

image copy

Copies the entire image in the active partition to the standby partition, a mirror image.

Syntax	`image copy active-to-standby`
Parameters	`active-to-standby` — Enter to copy the entire image in the active partition to the standby partition, a mirror image.
Default	Not configured
Command Mode	EXEC
Usage Information	Duplicate the active, running software image to the standby image location.
Example	`OS10# image copy active-to-standby`
Supported Releases	10.2.0E or later

image download

Downloads a new software image to the local file system.

Syntax	`image download` *`file-url`*
Parameters	*`file-url`* — Set the path to the image file: • `ftp://userid:passwd@hostip:/filepath` — Enter the path to copy from the remote FTP server. • `http[s]://hostip:/filepath` — Enter the path to copy from the remote HTTP or HTTPS server.

- `scp://userid:passwd@hostip:/filepath` — Enter the path to copy from the remote SCP file system.
- `sftp://userid:passwd@hostip:/filepath` — Enter the path to copy from the remote SFTP file system.
- `tftp://hostip:/filepath` — Enter the path to copy from the remote TFTP file system.
- `usb://filepath` — Enter the path to copy from the USB file system.

Default Not configured

Command Mode EXEC

Usage Information Use the `show image status` command to view the progress.

Example

```
OS10# image download ftp://admin@10.206.28.174:/PKGS_OS10-Enterprise-10.3.2E.55-installer-x86_64.bin
```

```
OS10# image download ftp://admin@10.206.28.174:/PKGS_OS10-Enterprise-10.4.0E.55-installer-x86_64.bin
```

Supported Releases 10.2.0E or later

image install

Installs a new image, either from a previously downloaded file or from a remote location.

Syntax `image install file-url`

Parameters

- *file-url* — Location of the image file:
  - `ftp://userid:passwd@hostip:/filepath` — Enter the path to install from a remote FTP server.
  - `http[s]://hostip:/filepath` — Enter the path to install from the remote HTTP or HTTPS server.
  - `scp://userid:passwd@hostip:/filepath` — Enter the path to install from a remote SCP file system.
  - `sftp://userid:passwd@hostip:/filepath` — Enter the path to install from a remote SFTP file system.
  - `tftp://hostip:/filepath` — Enter the path to install from a remote TFTP file system.
  - `image://filename` — Enter the path to install from a local file system.
  - `usb://filepath` — Enter the path to install from the USB file system.

Default All

Command Mode EXEC

Usage Information Use the `show image status` command to view the installation progress.

Example

```
OS10# image install ftp://10.206.28.174:/PKGS_OS10-Enterprise-10.3.2E.55-installer-x86_64.bin
```

```
OS10# image install ftp://10.206.28.174:/PKGS_OS10-Enterprise-10.4.0E.55-installer-x86_64.bin
```

Supported Releases 10.2.0E or later

show boot

Displays boot partition-related information.

Syntax `show boot [detail]`

Parameters `detail` — (Optional) Enter to display detailed information.

Default Not configured

Command Mode EXEC

Usage Information Use the `boot system` command to set the boot partition for the next reboot.

Example

```
OS10# show boot
Current system image information:
==================================
Type      Boot Type      Active                   Standby                  Next-Boot
-------------------------------------------------------------------------------------
Node-id 1 Flash Boot     [A] 10.4.3E             [B] 10.4.3E              [A] active
```

Example (Detail)

```
OS10# show boot detail
Current system image information detail:
=========================================
Type:                      Node-id 1
Boot Type:                 Flash Boot
Active Partition:          A
Active SW Version:         10.4.3E
Active SW Build Version:   10.4.3E.85
Active Kernel Version:     Linux 4.9.110
Active Build Date/Time:    2019-02-18T09:06:10Z
Standby Partition:         B
Standby SW Version:        10.4.3E
Standby SW Build Version:  10.4.3E.80
Standby Build Date/Time:   2019-02-17T15:36:08Z
Next-Boot:                 active[A]
```

Supported Releases 10.2.0E or later

show image status

Displays image transfer and installation information.

Syntax `show image status`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show image status
Image Upgrade State:   idle
================================================
File Transfer State:   idle
------------------------------------------------
  State Detail:        No download information available
  Task Start:          0000-00-00T00:00:00Z
  Task End:            0000-00-00T00:00:00Z
  Transfer Progress:   0 %
  Transfer Bytes:      0 bytes
```

```
  File Size:            0 bytes
  Transfer Rate:        0 kbps

Installation State:     idle
-----------------------------------------------
  State Detail:         No install information available
  Task Start:           0000-00-00T00:00:00Z
  Task End:             0000-00-00T00:00:00Z
```

Supported Releases 10.2.0E or later

show version

Displays software version information.

Syntax `show version`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
TR2# show version
Dell EMC Networking OS10 Enterprise
Copyright (c) 1999-2019 by Dell Inc. All Rights Reserved.
OS Version: 10.4.3.0
Build Version: 10.4.3.85
Build Time: 2019-02-18T17:06:10-0800
System Type: S4048-ON
Architecture: x86_64
Up Time: 2 days 05:58:01
```

Supported Releases 10.2.0E or later

10

OpenFlow

Switches implement the control plane and data plane in the same hardware. Software-defined network (SDN) decouples the software (control plane) from the hardware (data plane). A centralized SDN controller handles the control plane traffic and hardware configuration for data plane flows.

The SDN controller is the "brain" of an SDN. The SDN controller uses north-bound application programming interfaces (APIs) to communicate with the business logic applications and south-bound APIs to set up controlled network devices, such as OS10 switches.

OpenFlow is an implementation of SDN. OpenFlow enables programmable networks. You can develop SDN controller network applications using representational state transfer (REST) or JAVA APIs (north-bound APIs) to business logic applications. The SDN controller uses OpenFlow south-bound APIs to communicate with the switches and relay information from business logic applications.

Advantages of an SDN include customization, accelerating new feature development, lower operating costs, and fostering an open, multi-vendor environment.

OS10 supports OpenFlow protocol versions 1.0 and 1.3.

OS10 supports OpenFlow-only mode. In this mode, the SDN controller controls data path of the switch. The OpenFlow pipeline processes all data packets.

NOTE: When the switch is in OpenFlow mode, all Layer 2 (L2) and Layer 3 (L3) protocols are disabled. Link-level protocols such as Link Layer Discovery Protocol (LLDP), Dot1x, and Virtual Link Trunking (VLT) are disabled as well.

NOTE: OpenFlow Hybrid mode is not supported.

Supported Platforms

- S4048–ON
- S4048T-ON
- S4100-ON
- S4248FB-ON
- S4248FBL-ON
- S6010–ON
- Z9100–ON
- Z9264F-ON

NOTE: S5148F-ON and S3048-ON are not supported.

OS10 OpenFlow implementation reserves VLANs 1 and 4095.

The following is a known OpenFlow restriction in OS10:

Converting the switch from OpenFlow mode back to Normal mode removes all OpenFlow configurations. The switch returns to the pre-Openflow status. The management, interface (maximum transmission unit (MTU) and LLDP), and authentication, authorization, and accounting (AAA) settings specified in the Normal mode are retained.

To start up the switch in Factory Default mode, you must:

1 Delete the startup configuration using the `delete startup-configuration` command.
2 Enter the `reload` command.

NOTE: Do not use the `no openflow` or `no mode openflow-only` command.

```
OS10# delete startup-configuration
OS10# reload
```

OpenFlow logical switch instance

In OpenFlow-only mode, you can configure only one logical switch instance. After you enable OpenFlow mode, create a logical switch instance. The logical switch instance is disabled by default. When the logical switch instance is enabled, the OpenFlow application starts the connection with the configured controller.

When you create an OpenFlow logical switch instance, all the physical interfaces are automatically added to it.

OpenFlow controller

OS10 is qualified with the following SDN controllers:

- RYU
- Open Network Operating System (ONOS)

To establish a connection with the controller, configure the IPv4 address of the controller and port ID in the OpenFlow logical switch instance. The default port is 6653. You can connect controllers to the switch in OOB Connection mode. However, you can use any of the front-panel ports as the management interface using the `in-band` command. The inband port is removed from the OpenFlow switch instance and is not controlled by the controller.

The management port MTU is 1532 and the inband port MTU is 9216.

OpenFlow uses the Transmission Control Protocol (TCP) and Transport Layer Security (TLS) protocol for communication.

If the OpenFlow switch loses connection with the controller, the switch immediately enters Fail Secure mode. All the flows the controller installs are retained on the switch. The flow entries are removed based on the hard or idle timeout that you configure.

OpenFlow version 1.3

This section provides information about OpenFlow version 1.3 specifications for OS10.

Ports

An OpenFlow switch supports the following OpenFlow ports:

Table 35. Supported port types

Port types	Support
Physical ports	Supported
Logical ports	Not supported
Reserved ports	
(Required) ALL	Supported
(Required) CONTROLLER	Supported
(Required) TABLE	Not supported
(Required) IN PORT	Not supported

Port types	Support
(Required) ANY	Supported
(Optional) LOCAL	Not supported
(Optional) NORMAL	Not supported
(Optional) FLOOD	Not supported

Flow table

An OpenFlow flow table consists of flow entries. Each flow table entry contains the following fields:

Table 36. Supported fields

Fields	Support
match_fields	Supported
priority	Supported
counters	Supported
instructions	Supported
timeouts	Supported
cookie	Not supported

Group table

Not supported

Meter table

Not supported

Instructions

Each flow entry contains a set of instructions that execute when a packet matches the entry.

Table 37. Supported instructions

Instruction	Support
(Optional) Meter meter id	Not supported
(Optional) Apply-Actions action(s)	Supported
(Optional) Clear-Actions	Not supported
(Required) Write-Actions action(s)	Supported
(Optional) Write-Metadata metadata/mask	Not supported
(Required) Goto-table next-table-id	Not supported

Action set

An action set associates with each packet.

Table 38. Supported action sets

Action set	Support
copy TTL inwards	Not supported
pop	Not supported
push-MPLS	Not supported
push-VLAN	Not supported
copy TTL outwards	Not supported
decrement TTL	Not supported
set	Supported (selective fields)
qos	Not supported
group	Not supported
output	Supported

Action types

An action type associates with each packet.

Table 39. Supported action types

Action type	Support
Output	Supported
Set-queue	Not supported
Drop	Supported
Group	Not supported
Push-tag/Pop-tag	Not supported
Set-field	Partially supported • Source MAC—Supported • Destination MAC—Supported • VLAN ID—Supported • VLAN PCP—Supported • IP DSCP—Supported
change-TTL	Not supported

Counters

Counters are used for statistical purposes.

Table 40. Supported counters

Required/Optional	Counter	Bits	Support
Per flow table			
Required	Reference count (active entries)	32	Supported
Optional	Packet lookups	64	Supported
Optional	Packet matches	64	Supported
Per flow entry			
Optional	Received packets	64	Supported
Optional	Received bytes	64	Supported
Required	Duration (seconds)	32	Supported
Optional	Duration (nanoseconds)	32	Supported
Per port			
Required	Received packets	64	Supported
Required	Transmitted packets	64	Supported
Optional	Received bytes	64	Supported
Optional	Transmitted bytes	64	Supported
Optional	Receive drops	64	Not supported
Optional	Transmit drops	64	Not supported
Optional	Receive errors	64	Supported
Optional	Transmit errors	64	Supported
Optional	Receive frame alignment errors	64	Not supported
Optional	Receive overrun errors	64	Not supported
Optional	Receive CRC errors	64	Supported
Optional	Collisions	64	Supported
Required	Duration (seconds)	32	Not supported
Optional	Duration (nanoseconds)	32	Not supported
Per queue			
Required	Transmit packets	64	Not supported
Optional	Transmit bytes	64	Not supported
Optional	Transmit overrun errors	64	Not supported
Required	Duration (seconds)	32	Not supported
Optional	Duration (nanoseconds)	32	Not supported
Per group			
Optional	Reference count (flow entries)	32	Not supported

Required/Optional	Counter	Bits	Support
Optional	Packet count	64	Not supported
Optional	Byte count	64	Not supported
Required	Duration (seconds)	32	Not supported
Optional	Duration (nanoseconds)	32	Not supported
Per group bucket			
Optional	Packet count	64	Not supported
Optional	Byte count	64	Not supported
Per meter			
Optional	Flow count	32	Not supported
Optional	Input packet count	64	Not supported
Optional	Input byte count	64	Not supported
Required	Duration (seconds)	32	Not supported
Optional	Duration (nanoseconds)	32	Not supported
Per meter band			
Optional	In-band packet count	64	Not supported
Optional	In-band byte count	64	Not supported

OpenFlow protocol

The OpenFlow protocol supports three message types, each with multiple subtypes:

- Controller-to-switch
- Asynchronous
- Symmetric

Controller-to-switch

Table 41. Supported controller-to-switch types

Controller-to-switch types	Supported/Not supported
Feature request	Supported
Configuration get	Supported
Configuration set	Supported
Modify-state	Supported
Read-state	Supported
Packet-out	Supported
Barrier	Supported
Role-request	Supported

Asynchronous

Table 42. Supported asynchronous types

Asynchronous types	Supported/Not supported
Packet-in	Supported
Flow-removed	Supported
Port-status	Supported
Error	Supported

Symmetric

Table 43. Supported symmetric types

Symmetric types	Supported/Not supported
Hello	Supported
Echo	Supported
Experimenter	Not supported

Connection setup TCP

Table 44. Supported modes

Modes	Supported/Not supported
Connection interruption	• fail-secure-mode—Supported • fail-standalone-mode—Not supported
TLS encryption	Supported
Multiple controller	Not supported
Auxiliary connections	Not supported
Number of logical switches	One

Supported controllers

REST APIs on

- RYU
- ONOS

Flow table modification messages

Table 45. Supported messages

Flow table modification messages	Supported/Not supported
OFPFC_ADD=0	Supported
OFPFC_MODIFY=1	Supported

Flow table modification messages	Supported/Not supported
OFPFC_MODIFY_STRICT=2	Supported
OFPFC_DELETE=3	Supported
OFCPC_DELETE_STRICT=4	Supported

Message types

Table 46. Supported message types

Message Type	Message	Support
Immutable messages	OFPT_HELLO=0	Supported
	OFPT_ERROR=1	Supported
	OFPT_ECHO_REQUEST=2	Supported
	OFPT_ECHO_REPLY=3	Supported
Switch configuration messages	OFPT_FEATURES_REQUEST=5	Supported
	OFPT_FEATURES_REPLY=6	Supported
	OFPT_GET_CONFIG_REQUEST=7	Supported
	OFPT_GET_CONFIG_REPLY=8	Supported
	OFPT_SET_CONFIG=9	Supported
Asynchronous messages	OFPT_PACKET_IN=10	Supported
	OFPT_FLOW_REMOVED=11	Supported
	OFPT_PORT_STATUS=12	Supported
Controller command messages	OFPT_PACKET_OUT=13	Supported
	OFPT_FLOW_MOD=14	Supported
	OFPT_GROUP_MOD=15	Not supported
	OFPT_PORT_MOD=16	Supported
	OFPT_TABLE_MOD=17	Not supported
Multipart messages	OFPT_MULTIPART_REQUEST=18	Supported
	OFPT_MULTIPART_REPLY=19	Supported
Barrier messages	OFPT_BARRIER_REQUEST=20	Supported
	OFPT_BARRIER_REPLY=21	Supported
Queue configuration messages	OFPT_QUEUE_GET_CONFIG_REQUEST=22	Not supported
	OFPT_QUEUE_GET_CONFIG_REPLY=23	Not supported
Controller role change request messages	OFPT_ROLE_REQUEST=24	Not supported
	OFPT_ROLE_REPLY=25	Not supported
Asynchronous message configuration	OFPT_GET_ASYNC_REQUEST=26	Not supported
	OFPT_GET_ASYNC_REPLY=27	Not supported

Message Type	Message	Support
	OFPT_SET_ASYNC=28	Not supported
Meters and rate limiters configuration messages	OFPT_METER_MOD=29	Not supported

Flow match fields

Table 47. Supported fields

Flow match fields	Supported/Not supported
OFPXMT_OFB_IN_PORT = 0	Supported
OFPXMT_OFB_IN_PHY_PORT = 1	Not supported
OFPXMT_OFB_METADATA = 2	Not supported
OFPXMT_OFB_ETH_DST = 3	Supported
OFPXMT_OFB_ETH_SRC = 4	Supported
OFPXMT_OFB_ETH_TYPE = 5	Supported
OFPXMT_OFB_VLAN_VID = 6	Supported
OFPXMT_OFB_VLAN_PCP = 7	Supported
OFPXMT_OFB_IP_DSCP = 8	Supported
OFPXMT_OFB_IP_ECN = 9	Supported
OFPXMT_OFB_IP_PROTO = 10	Supported
OFPXMT_OFB_IPV4_SRC = 11	Supported
OFPXMT_OFB_IPV4_DST = 12	Supported
OFPXMT_OFB_TCP_SRC = 13	Supported
OFPXMT_OFB_TCP_DST = 14	Supported
OFPXMT_OFB_UDP_SRC = 15	Supported
OFPXMT_OFB_UDP_DST = 16	Supported
OFPXMT_OFB_SCTP_SRC = 17	Not supported
OFPXMT_OFB_SCTP_DST = 18	Not supported
OFPXMT_OFB_ICMPV4_TYPE = 19	Supported
OFPXMT_OFB_ICMPV4_CODE = 20	Supported

Flow match fields	Supported/Not supported
OFPXMT_OFB_ARP_OP = 21	Not supported
OFPXMT_OFB_ARP_SPA = 22	Not supported
OFPXMT_OFB_ARP_TPA = 23	Not supported
OFPXMT_OFB_ARP_SHA = 24	Not supported
OFPXMT_OFB_ARP_THA = 25	Not supported
OFPXMT_OFB_IPV6_SRC = 26	Not supported
OFPXMT_OFB_IPV6_DST = 27	Not supported
OFPXMT_OFB_IPV6_FLABEL = 28	Not supported
OFPXMT_OFB_ICMPV6_TYPE = 29	Not supported
OFPXMT_OFB_ICMPV6_CODE = 30	Not supported
OFPXMT_OFB_IPV6_ND_TARGET = 31	Not supported
OFPXMT_OFB_IPV6_ND_SLL = 32	Not supported
OFPXMT_OFB_IPV6_ND_TLL = 33	Not supported
OFPXMT_OFB_MPLS_LABEL = 34	Not supported
OFPXMT_OFB_MPLS_TC = 35	Not supported
OFPXMT_OFP_MPLS_BOS = 36	Not supported
OFPXMT_OFB_PBB_ISID = 37	Not supported
OFPXMT_OFB_TUNNEL_ID = 38	Not supported
OFPXMT_OFB_IPV6_EXTHDR = 39	Not supported

Action structures

Table 48. Supported action structures

Action structures	Supported/Not supported
OFPAT_OUTPUT = 0	Supported
OFPAT_COPY_TTL_OUT = 11	Not supported
OFPAT_COPY_TTL_IN = 12	Not supported
OFPAT_SET_MPLS_TTL = 15	Not supported
OFPAT_DEC_MPLS_TTL = 16	Not supported

Action structures	Supported/Not supported
OFPAT_PUSH_VLAN = 17	Not supported
OFPAT_POP_VLAN = 18	Not supported
OFPAT_PUSH_MPLS = 19	Not supported
OFPAT_POP_MPLS = 20	Not supported
OFPAT_SET_QUEUE = 21	Not supported
OFPAT_GROUP = 22	Not supported
OFPAT_SET_NW_TTL = 23	Not supported
OFPAT_DEC_NW_TTL = 24	Not supported
OFPAT_SET_FIELD = 25	Supported
OFPAT_PUSH_PBB = 26	Not supported
OFPAT_POP_PBB = 27	Not supported

Capabilities supported by the data path

Table 49. Supported capabilities

Capabilities	Supported/Not supported
OFPC_FLOW_STATS = 1 << 0	Supported
OFPC_TABLE_STATS = 1 << 1	Not supported
OFPC_PORT_STATS = 1 << 2	Supported
OFPC_GROUP_STATS = 1 << 3	Not supported
OFPC_IP_REASM = 1 << 5	Not supported
OFPC_QUEUE_STATS = 1 << 6	Not supported
OFPC_PORT_BLOCKED = 1 << 8	Not supported

Multipart message types

Table 50. Supported message types

Message type description	Request/Reply Body	Message	Support
Description of this OpenFlow switch	• The request body is empty	OFPMP_DESC = 0	Supported

Message type description	Request/Reply Body	Message	Support
	• The reply body is struct ofp_desc		
Individual flow statistics	• The request body is struct ofp_flow_stats_request • The reply body is an array of struct ofp_flow_stats	OFPMP_FLOW = 1	Supported
Aggregate flow statistics	• The request body is struct ofp_aggregate_stats_request • The reply body is struct ofp_aggregate_stats_reply	OFPMP_AGGREGATE = 2	Supported
Flow table statistics	• The request body is empty • The reply body is an array of struct ofp_table_stats	OFPMP_TABLE = 3	Supported
Port statistics	• The request body is struct ofp_port_stats_request • The reply body is an array of struct ofp_port_stats	OFPMP_PORT_STATS = 4	Supported
Queue statistics for a port	• The request body is struct ofp_queue_stats_request • The reply body is an array of struct ofp_queue_stats	OFPMP_QUEUE = 5	Not supported
Group counter statistics	• The request body is struct ofp_group_stats_request • The reply is an array of struct ofp_group_stats	OFPMP_GROUP = 6	Not supported
Group description	• The request body is empty • The reply body is an array of struct ofp_group_desc_stats	OFPMP_GROUP_DESC = 7	Not supported
Group features	• The request body is empty • The reply body is struct ofp_group_features	OFPMP_GROUP_FEATURES = 8	Not supported
Meter statistics	• The request body is struct ofp_meter_multipart_requests • The reply body is an array of struct ofp_meter_stats	OFPMP_METER = 9	Not supported
Meter configuration	• The request body is struct ofp_meter_multipart_requests • The reply body is an array of struct ofp_meter_config	OFPMP_METER_CONFIG = 10	Not supported
Meter features	• The request body is empty	OFPMP_METER_FEATURES = 11	Not supported

Message type description	Request/Reply Body	Message	Support
	• The reply body is struct ofp_meter_features		
Table features	• The request body is empty or contains an array of struct ofp_table_features that includes the controller's desired view of the switch. If the switch is unable to set the specified view an error is returned • The reply body is an array of struct ofp_table_features	OFPMP_TABLE_FEATURES = 12	Supported
Port description	• The request body is empty • The reply body is an array of struct ofp_port	OFPMP_PORT_DESC = 13	Supported

Switch description

The OFPMP_DESC multipart request type includes information about the switch manufacturer, hardware revision, software revision, serial number, and description.

Table 51. Supported descriptions

Switch description	Supported/Not supported
char mfr_desc[DESC_STR_LEN]	Supported
char hw_desc[DESC_STR_LEN]	Supported
char sw_desc[DESC_STR_LEN]	Supported
char serial_num[SERIAL_NUM_LEN]	Supported
char dp_desc[DESC_STR_LEN]	Supported

Property type

Table 52. Supported properties

Property type	Supported/Not supported
OFPTFPT_INSTRUCTIONS = 0	Supported
OFPTFPT_INSTRUCTIONS_MISS = 1	Not supported
OFPTFPT_NEXT_TABLES = 2	Not supported
OFPTFPT_NEXT_TABLES_MISS = 3	Not supported
OFPTFPT_WRITE_ACTIONS = 4	Supported
OFPTFPT_WRITE_ACTIONS_MISS = 5	Not supported
OFPTFPT_APPLY_ACTIONS = 6	Supported

Property type	Supported/Not supported
OFPTFPT_APPLY_ACTIONS_MISS = 7	Not supported
OFPTFPT_MATCH = 8	Supported
OFPTFPT_WILDCARDS = 10	Supported
OFPTFPT_WRITE_SETFIELD = 12	Supported
OFPTFPT_WRITE_SETFIELD_MISS = 13	Not supported
OFPTFPT_APPLY_SETFIELD = 14	Supported
OFPTFPT_APPLY_SETFIELD_MISS = 15	Not supported

Group configuration

Table 53. Supported configurations

Group configuration	Supported/Not supported
OFPGFC_SELECT_WEIGHT = 1 << 0	Not supported
OFPGFC_SELECT_LIVENESS = 1 << 1	Not supported
OFPGFC_CHAINING = 1 << 2	Not supported
OFPGFC_CHAINING_CHECKS = 1 << 3	Not supported

Controller roles

Table 54. Supported controller roles

Controller roles	Supported/Not supported
OFPCR_ROLE_NOCHANGE = 0	Not supported
OFPCR_ROLE_EQUAL = 1	Supported
OFPCR_ROLE_MASTER = 2	Supported
OFPCR_ROLE_SLAVE = 3	Not supported

Packet-in reasons

Table 55. Supported reasons

Packet-in reasons	Supported/Not supported
OFPR_NO_MATCH = 0	Supported
OFPR_ACTION = 1	Supported
OFPR_INVALID_TTL = 2	Not supported

Flow-removed reasons

Table 56. Supported reasons

Flow-removed reasons	Supported/Not supported
OFPRR_IDLE_TIMEOUT = 0	Supported
OFPRR_HARD_TIMEOUT = 1	Supported
OFPRR_DELETE = 2	Supported
OFPRR_GROUP_DELETE = 3	Not supported

Error types from switch to controller

Table 57. Supported error types

Error types	Supported/Not supported
OFPET_HELLO_FAILED = 0	Supported
OFPET_BAD_REQUEST = 1	Supported
OFPET_BAD_ACTION = 2	Supported
OFPET_BAD_INSTRUCTION = 3	Supported
OFPET_BAD_MATCH = 4	Supported
OFPET_FLOW_MOD_FAILED = 5	Supported
OFPET_GROUP_MOD_FAILED = 6	Not supported
OFPET_PORT_MOD_FAILED = 7	Supported
OFPET_TABLE_MOD_FAILED = 8	Not supported
OFPET_QUEUE_OP_FAILED = 9	Not supported
OFPET_SWITCH_CONFIG_FAILED = 10	Not supported
OFPET_ROLE_REQUEST_FAILED = 11	Not supported
OFPET_METER_MOD_FAILED = 12	Not supported
OFPET_TABLE_FEATURES_FAILED = 13	Not supported
Bad request code	
OFPBRC_BAD_VERSION = 0	Supported

Error types	Supported/Not supported
OFPBRC_BAD_TYPE = 1	Supported
OFPBRC_BAD_MULTIPART = 2	Not supported
OFPBRC_BAD_EXPERIMENTER = 3	Not supported
OFPBRC_BAD_EXP_TYPE = 4	Not supported
OFPBRC_EPERM = 5	Not supported
OFPBRC_BAD_LEN = 6	Supported
OFPBRC_BUFFER_EMPTY = 7	Not supported
OFPBRC_BUFFER_UNKNOWN = 8	Not supported
OFPBRC_BAD_TABLE_ID = 9	Supported
OFPBRC_IS_SLAVE = 10	Not supported
OFPBRC_BAD_PORT = 11	Supported
OFPBRC_BAD_PACKET = 12	Not supported
OFPBRC_MULTIPART_BUFFER_OVERFLOW = 13	Not supported
Bad action code	
OFPBAC_BAD_TYPE = 0	Supported
OFPBAC_BAD_LEN = 1	Supported
OFPBAC_BAD_EXPERIMENTER = 2	Not supported
OFPBAC_BAD_EXP_TYPE = 3	Not supported
OFPBAC_BAD_OUT_PORT = 4	Supported
OFPBAC_BAD_ARGUMENT = 5	Supported
OFPBAC_EPERM = 6	Not supported
OFPBAC_TOO_MANY = 7	Supported
OFPBAC_BAD_QUEUE = 8	Not supported
OFPBAC_BAD_OUT_GROUP = 9	Not supported
OFPBAC_MATCH_INCONSISTENT = 10	Not supported
OFPBAC_UNSUPPORTED_ORDER = 11	Not supported
OFPBAC_BAD_TAG = 12	Not supported

Error types	Supported/Not supported
OFPBAC_BAD_SET_TYPE = 13	Not supported
OFPBAC_BAD_SET_LEN = 14	Not supported
OFPBAC_BAD_SET_ARGUMENT = 15	Supported
Bad instruction code	
OFPBIC_UNKNOWN_INST = 0	Not supported
OFPBIC_UNSUP_INST = 1	Not supported
OFPBIC_BAD_TABLE_ID = 2	Not supported
OFPBIC_UNSUP_METADATA = 3	Not supported
OFPBIC_UNSUP_METADATA_MASK = 4	Not supported
OFPBIC_BAD_EXPERIMENTER = 5	Not supported
OFPBIC_BAD_EXP_TYPE = 6	Not supported
OFPBIC_BAD_LEN = 7	Not supported
OFPBIC_EPERM = 8	Not supported
Bad match code	
OFPBMC_BAD_TYPE = 0	Not supported
OFPBMC_BAD_LEN = 1	Not supported
OFPBMC_BAD_TAG = 2	Not supported
OFPBMC_BAD_DL_ADDR_MASK = 3	Not supported
OFPBMC_BAD_NW_ADDR_MASK = 4	Not supported
OFPBMC_BAD_WILDCARDS = 5	Not supported
OFPBMC_BAD_FIELD = 6	Not supported
OFPBMC_BAD_VALUE = 7	Not supported
OFPBMC_BAD_MASK = 8	Not supported
OFPBMC_BAD_PREREQ = 9	Not supported
OFPBMC_DUP_FIELD = 10	Not supported
OFPBMC_EPERM = 11	Not supported
Flow modification failed code	

Error types	Supported/Not supported
OFPFMFC_UNKNOWN = 0	Supported
OFPFMFC_TABLE_FULL = 1	Supported
OFPFMFC_BAD_TABLE_ID = 2	Supported
OFPFMFC_OVERLAP = 3	Supported
OFPFMFC_EPERM = 4	Not supported
OFPFMFC_BAD_TIMEOUT = 5	Not supported
OFPFMFC_BAD_COMMAND = 6	Supported
OFPFMFC_BAD_FLAGS = 7	Not supported
Group modification failed code	
OFPGMFC_GROUP_EXISTS = 0	Not supported
OFPGMFC_INVALID_GROUP = 1	Not supported
OFPGMFC_WEIGHT_UNSUPPORTED = 2	Not supported
OFPGMFC_OUT_OF_GROUPS = 3	Not supported
OFPGMFC_OUT_OF_BUCKETS = 4	Not supported
OFPGMFC_CHAINING_UNSUPPORTED = 5	Not supported
OFPGMFC_WATCH_UNSUPPORTED = 6	Not supported
OFPGMFC_LOOP = 7	Not supported
OFPGMFC_UNKNOWN_GROUP = 8	Not supported
OFPGMFC_CHAINED_GROUP = 9	Not supported
OFPGMFC_BAD_TYPE = 10	Not supported
OFPGMFC_BAD_COMMAND = 11	Not supported
OFPGMFC_BAD_BUCKET = 12	Not supported
OFPGMFC_BAD_WATCH = 13	Not supported
OFPGMFC_EPERM = 14	Not supported
Port modification failed code	
OFPPMFC_BAD_PORT = 0	Supported
OFPPMFC_BAD_HW_ADDR = 1	Supported

Error types	Supported/Not supported
OFPPMFC_BAD_CONFIG = 2	Not supported
OFPPMFC_BAD_ADVERTISE = 3	Not supported
OFPPMFC_EPERM = 4	Not supported
Table modification failed code	
OFPTMFC_BAD_TABLE = 0	Supported
OFPTMFC_BAD_CONFIG = 1	Not supported
OFPTMFC_EPERM = 2	Not supported
Queue operation failed code	
OFPQOFC_BAD_PORT = 0	Supported
OFPQOFC_BAD_QUEUE = 1	Not supported
OFPQOFC_EPERM = 2	Not supported
Switch configuration failed code	
OFPSCFC_BAD_FLAGS = 0	Not supported
OFPSCFC_BAD_LEN = 1	Not supported
OFPSCFC_EPERM = 2	Not supported
Role request failed code	
OFPRRFC_STALE = 0	Not supported
OFPRRFC_UNSUP = 1	Not supported
OFPRRFC_BAD_ROLE = 2	Not supported
Table features failed code	
OFPTFFC_BAD_TABLE = 0	Supported
OFPTFFC_BAD_METADATA = 1	Not supported
OFPTFFC_BAD_TYPE = 2	Not supported
OFPTFFC_BAD_LEN = 3	Not supported
OFPTFFC_BAD_ARGUMENT = 4	Not supported
OFPTFFC_EPERM = 5	Not supported

OpenFlow use cases

OS10 OpenFlow protocol support allows the flexibility of using vendor-neutral applications and to use applications that you create. For example, the OS10 OpenFlow implementation supports L2 applications similar to the ones found in the following websites:

- https://github.com/osrg/ryu/tree/master/ryu/app (only L2 applications are supported)
- https://github.com/osrg/ryu/tree/master/ryu/app

NOTE: OS10 supports applications based on OpenFlow versions 1.0 and 1.3.

- **Switching loop removal**

  Consider the case of a single broadcast domain where switching loops are common. This issue occurs because of redundant paths in an L2 network.

  Switching loops create broadcast storms with broadcasts and multicasts being forwarded out of every switch port. Every switch in the network repeatedly re-broadcasts the messages and floods the entire network.

  To solve broadcast storms in an OpenFlow network, a centralized controller makes all the control plane decisions and manages the switches. The controller has the complete view of the topology. MAC address learning is centralized. OpenFlow identifies the correct path and forwards the packets to the relevant switch thereby avoiding switching loops.

- **Reactive flow installation**

  Consider the case of dynamic learning of flows for bidirectional traffic. Flows are learnt as and when a packet arrives.

  With dynamic learning in an OpenFlow network, the OpenFlow switch receives a packet that does not match the flow table entries and sends the packet to the SDN controller to process it. The controller identifies the path the packet has to traverse and updates the flow table with a new entry. The controller also decides the caching time of the flow table entries.

Configure OpenFlow

Ensure IP connectivity between the switch and the controller. When you convert the switch from Normal mode to OpenFlow mode, the switch retains the management, interface, and AAA settings.

The following example lists the minimum configuration needed to establish the connection between the OpenFlow controller and a logical switch instance.

1 Enter the OPENFLOW configuration mode.

```
OS10# configure terminal
OS10 (config)# openflow
OS10 (config-openflow)#
```

2 Enable the OpenFlow-only mode.

```
OS10 (config-openflow)# mode openflow-only
```

The system prompts you to reload the switch. Enter `yes` to enable OpenFlow-only mode.

NOTE: When the switch starts up in OpenFlow mode, it disables all L2 and 3 protocols. Many CLI commands are not available when the switch is in OpenFlow-only mode. For a list of commands that are available in OpenFlow-only mode, see CLI commands available in the OpenFlow-only mode.

3 Configure a logical switch instance.

a Option 1; for OOB management:

1 Configure an IP address for the management port. Ensure that there is IP connectivity between the switch and the controller.

```
OS10# configure terminal
OS10 (config)# interface management 1/1/1
OS10 (conf-if-ma-1/1/1)# ip address 11.1.1.1/24
OS10 (conf-if-ma-1/1/1)# no shutdown
OS10 (conf-if-ma-1/1/1)# exit
```

2 Configure the logical switch instance, *of-switch-1*.

```
OS10# configure terminal
OS10 (config)# openflow
OS10 (config-openflow)# switch of-switch-1
```

b Option 2; for in-band management:

1 Configure one of the front-panel ports as the management port.

```
OS10# configure terminal
OS10 (config)# openflow
OS10 (config-openflow)# in-band-mgmt interface ethernet 1/1/1
OS10 (config-openflow)#
```

2 Configure an IPv4 address on the front-panel management port.

```
OS10# configure terminal
OS10 (config)# interface ethernet 1/1/1
OS10 (conf-if-eth1/1/1)# ip address 11.1.1.1/24
OS10 (conf-if-eth1/1/1)# no shutdown
```

3 Configure the logical switch instance, *of-switch-1*.

```
OS10# configure terminal
OS10 (config)# openflow
OS10 (config-openflow)# switch of-switch-1
```

4 Configure the OpenFlow controller to establish a connection with the logical switch instance.

```
OS10 (config-openflow-switch)# controller ipv4 ip-address port port-id
```

```
OS10 (config-openflow-switch)# controller ipv4 10.1.1.1 port 6633
```

where `a.b.c.d` is the IP address of the controller and port 6633 is for OpenFlow communication.

5 Enter the `no shutdown` command to enable the logical switch instance.

```
OS10 (config-openflow-switch) no shutdown
```

Establish TLS connection

- Generate the switch and controller certificates from a server that supports public-key infrastructure (PKI). You need the following certificates:
  - Controller certificate
  - Switch certificate
  - Private key file to verify the switch certificate
- The certificates and private key files must be in the Privacy-Enhanced Mail (PEM) format.

For certificate-based authentication, you must establish a TLS connection between the switch and the controller before you configure OpenFlow on the switch. The following procedure explains how to install the controller and switch certificates on the OS10 switch. Refer to the controller documentation for information on how to install the certificates on the controller.

NOTE: This procedure is optional. Use this procedure if you want to configure certificate-based authentication between the switch and the controller.

1 Log in to the OS10 switch with administrator credentials.

2 Enter the following command to copy the certificates to the OS10 switch.

In the following commands, the destination path and the destination file name on the OS10 switch, for example, `config://../openflow/cacert.pem`, remain the same in your deployment. Ensure that you enter the destination path and destination file names as specified in the following example:

```
OS10# copy scp://username:password@server-ip/full-path-to-the-certificates/controller-cert.pem config://../openflow/cacert.pem
OS10# copy scp://username:password@server-ip/full-path-to-the-certificates/switch-cert.pem config://../openflow/sc-cert.pem
OS10# copy scp://username:password@server-ip/full-path-to-the-certificates/switch-privkey.pem config://../openflow/sc-privkey.pem
```

where *server-ip* refers to the server where you have stored the certificates, and *username* and *password* refers to the credentials you need to access the server with the certificates.

3 Perform the steps described in the Configure OpenFlow protocol on the switch topic to configure OpenFlow.

OpenFlow commands

controller

Configures an OpenFlow controller that the logical switch instance connects to.

Syntax `controller ipv4 ipv4-address [port port-number] [security {none|tls}]`

Parameters

- `ipv4 ipv4-address`—Enter `ipv4`, then the IP address of the controller.
- `port port-number`—Enter the keyword, then the port number, from 1 to 65,535. The default port is 6653.
- `security {none|tls}`—Specify the type of connection. The default is `security none`. The TCP connection is used.

Default TCP. The default port number is 6653.

Command Mode OPENFLOW SWITCH CONFIGURATION

Usage Information If you specify the `security tls` option, the OpenFlow application looks for the following certificates and private key in the following locations specified for certificate-based authentication. For information about obtaining certificates and installing them on the switch and the controller, see Establish TLS connection between the switch and the controller.

ca_cert (certificate that identifies the controller as being trustworthy)	`/config/etc/opt/dell/os10/openflow/cacert.pem`
certificate (certificate that identifies the switch as being trustworthy)	`/config/etc/opt/dell/os10/openflow/sc-cert.pem`
private key (the private key corresponding to the switch certificate)	`/config/etc/opt/dell/os10/openflow/sc-privkey.pem`

Example The following example configures an OpenFlow controller with IP address 10.11.63.56 on port 6633 for the logical switch instance, of-switch-1.

```
OS10# configure terminal
OS10 (config)# openflow
OS10 (config-openflow)# switch of-switch-1
OS10 (config-openflow-switch)# controller ipv4 10.11.63.56 port 6633
OS10 (config-openflow-switch)#
```

Supported Releases 10.4.1 or later

dpid-mac-address

Specifies the MAC address bits of the datapath ID (DPID) of the logical switch instance.

Syntax `dpid-mac-address MAC-address`

Parameters *MAC-address*—48-bit MAC address in hexadecimal notation, nn:nn:nn:nn:nn:nn

Default MAC address

Command Mode OPENFLOW SWITCH CONFIGURATION

Usage Information The controller uses the DPID to identify the logical switch instance. The DPID is a 64-bit number that is sent to the controller in the `features_reply` message. The DPID is constructed from the instance ID, which is the most significant 16 bits (default to 0) and the DPID-MAC-ADDRESS, which is the least significant 48 bits. OS10 currently supports only one logical switch instance and the instance ID is automatically set to 0. This value is not configurable.

You can use this command to modify the MAC address bits of the DPID.

Example DPID MAC address is 00:00:00:00:00:0a.

```
OS10 (config)# openflow
OS10 (config-openflow)# switch of-switch-1
OS10 (config-openflow-switch)# dpid-mac-address 00:00:00:00:00:0a
OS10 (config-openflow-switch)#
```

Supported Releases 10.4.1 or later

in-band-mgmt

Configures the front-panel ports as the management interface that the SDN controller connects to.

Syntax `in-band-mgmt interface ethernet node/slot/port[:subport]`

Parameters *node/slot/port[:subport]*—Enter the physical port information.

Default None

Command Mode OPENFLOW CONFIGURATION

Usage Information Use this command to convert any one of the front-panel ports as the management interface. This port is not part of the OpenFlow logical switch instance. All the ports are L2 ports by default. If you configure one of the front-panel ports as the management interface, the port becomes an L3 port. You can configure an L3 IPv4 address only to the front-panel port that you have specified in this command. Ensure that you have IP connectivity between the specified port and the controller.

The `no` form of this command removes this configuration and the front-panel port becomes part of the OpenFlow logical switch instance.

Example

```
OS10# configure terminal
OS10(config)# openflow
```

```
OS10 (config-openflow)# in-band-mgmt interface ethernet 1/1/1
OS10 (config-openflow)# no shutdown
```

Supported Releases 10.4.1 or later

max-backoff

Configures the time interval, in seconds, that the logical switch instance waits after requesting a connection with the OpenFlow controller.

Syntax `max-backoff interval`

Parameters *interval*—Enter the amount of time, in seconds, that the logical switch instance waits after it attempts to establish a connection with the OpenFlow controller, from 1 to 65,535.

Default 8 seconds

Command Mode OPENFLOW SWITCH CONFIGURATION

Usage Information If the interval time lapses, the logical switch instance re-attempts to establish a connection with the OpenFlow controller.

Example

```
OS10 (config)# openflow
OS10 (config-openflow)# switch of-switch-1
OS10 (config-openflow-switch)# max-backoff 25
OS10 (config-openflow-switch)#
```

Supported Releases 10.4.1 or later

mode openflow-only

Enables OpenFlow-only mode on the switch.

Syntax `mode openflow-only`

Parameters None

Default None

Command Mode OPENFLOW CONFIGURATION

Usage Information Use this command to enable OpenFlow-only mode. This command reloads the switch and boots to OpenFlow-only mode. This command deletes all L2 and L3 configurations. However, the system management and AAA configurations are retained.

The `no` form of this command prompts you to reload the switch. If you enter `yes`, the switch deletes all OpenFlow configurations, including the controller IP, port, certificates, and reloads, then returns to the Normal mode.

NOTE: For a list of available commands when the switch is in the OpenFlow-only mode, see CLI commands available in the OpenFlow-only mode.

Example

```
OS10 (config-openflow)# mode openflow-only
OS10 (config-openflow)#
```

Supported Releases 10.4.1 or later

openflow

Enters OPENFLOW configuration mode.

Syntax `openflow`

Parameters None

Default None

Command Mode CONFIGURATION

Usage Information All OpenFlow configurations are performed in this mode.

The `no` form of this command prompts a switch reload. If you enter `yes`, the system deletes all OpenFlow configurations and the switch returns to the normal mode after the reload.

Example

```
OS10# configure terminal
OS10(config)# openflow
OS10 (config-openflow)#
```

Supported Releases 10.4.1 or later

probe-interval

Configures the echo request interval, in seconds, for the controller configured with the logical switch instance.

Syntax `probe-interval` *`interval`*

Parameters *`interval`*—Enter the amount of time, in seconds, between the `keepalive` messages, also known as echo requests, from 1 to 65,535.

Default 5 seconds

Command Mode OPENFLOW SWITCH CONFIGURATION

Usage Information None

Example

```
OS10 (config)# openflow
OS10 (config-openflow)# switch of-switch-1
OS10 (config-openflow-switch)# probe-interval 20
OS10 (config-openflow-switch)#
```

Supported Releases 10.4.1 or later

protocol-version

Specifies protocol version the logical switch interface uses.

Syntax `protocol-version` *`version`*

Parameters *`version`*—Choose from one of the following:

- `negotiate`—Enter the keyword to negotiate versions 1.0 or 1.3 with the controller. The highest of the supported versions is selected.
- `1.0`—Specify the logical switch instance OpenFlow protocol version as 1.0.
- `1.3`—Specify the logical switch instance OpenFlow protocol version as 1.3.

Default `negotiate`

Command Mode OPENFLOW SWITCH CONFIGURATION

Usage Information

NOTE: Only use this command should be run when the logical switch instance is disabled. Use the `shutdown` command to disable the logical switch instance. After you run this command, enter the `no shutdown` command to enable the logical switch instance again.

- When you specify, `negotiate`, the switch negotiates versions 1.0 and 1.3 and selects the highest of the versions supported by the controller. The negotiation is based on the hello handshake described in the OpenFlow Specification 1.3.
- When you specify, `1.0`, the switch establishes a connection with the controller that supports version 1.0 only.
- When you specify, `1.3`, the switch establishes a connection with the controller that supports version 1.3 only.

Example The following example shows a logical switch instance, of-switch-1, configured to interact with controllers that support the OpenFlow protocol version 1.3.

```
OS10 (config)# openflow
OS10 (config-openflow)# switch of-switch-1
OS10 (config-openflow-switch)# shutdown
OS10 (config-openflow-switch)# protocol-version 1.3
OS10 (config-openflow-switch)# no shutdown
OS10 (config-openflow-switch)#
```

Supported Releases 10.4.1 or later

rate-limit packet_in

Configures the maximum packet rate for the controller connection, and the maximum packets permitted in a burst sent to the controller in a second.

Syntax `rate-limit packet_in` *`controller-packet-rate`* `[burst` *`maximum-packets-to-controller`*`]`

Parameters

- *`controller-packet-rate`*—Rate in packets per second for the controller OpenFlow channel connection, from 100 to 268000000 seconds. The default is 0 seconds, disabled.
- *`maximum-packets-to-controller`*—Burst in packets for the controller OpenFlow channel connection, from 25 to 1073000. The default is 0 seconds, disabled. This parameter is optional. It is set to 25% of the configured rate value, if not configured.

Default Disabled

Command Mode OPENFLOW SWITCH CONFIGURATION

Usage Information OpenFlow sets the specified rate and burst for the controller's connection with the logical switch instance. The actual rate and burst on the controller has a maximum of two times the configured values. For example, when you configure a rate of 1000 PPS and a burst of 300 packet bursts per second, the packets can egress on the connection at rates of up to 2000 PPS and 600 packet bursts per second.

The `no` form of this command disables rate limiting on the controller connection.

NOTE: This command is a software rate limiting command and applies only to the OpenFlow channel connection between the controller and the logical switch instance. This command is not related to the switch's data-plane rate limits.

Example

The following example configures a logical switch instance, of-switch-1, with an OpenFlow controller at a rate of 1000 PPS and packet bursts of 300 packets.

```
OS10 (config)# openflow
OS10 (config-openflow)# switch of-switch-1
OS10 (config-openflow-switch)# controller ipv4 10.11.63.56 port 6633
OS10 (config-openflow-switch)# rate-limit packet_in 1000 burst 300
OS10 (config-openflow-switch)#
```

Supported Releases 10.4.1 or later

show openflow

Displays general OpenFlow switch and the logical switch instance information.

Syntax `show openflow`

Parameters None

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show openflow

Manufacturer                 : DELL
Hardware Description         :
Software Description         : Dell Networking OS10-Premium, Dell Networking
Application Software Version: 10.4.9999EX
Serial Number                :
Capabilities                 : port, table, flow
Switch mode                  : openflow-only
Match fields                 :
    Layer-1 : in-port
    Layer-2 : eth-src, eth-dst, eth-type, vlan-id, vlan-pcp
    Layer-3 : ipv4-src, ipv4-dst, ip-protocol, ip-dscp, ip-ecn
    Layer-4 : tcp-src, tcp-dst, udp-src, udp-dst, icmpv4-type, icmpv4-code
Instructions                 : apply-actions, write-actions
Actions                      : output, set-field
Set field actions            : eth-src, eth-dst, vlan-id, vlan-pcp, ip-dscp
TLS parameters               :
    certificate identifying trustworthy controller : /config/etc/opt/dell/os10/
openflow/cacert.pem
    certificate identifying trustworthy switch     : /config/etc/opt/dell/os10/
openflow/sc-cert.pem
    private key                                    : /config/etc/opt/dell/os10/
openflow/sc-privkey.pem
```

Supported Releases 10.4.1 or later

show openflow flows

Displays OpenFlow flows for a specific logical switch instance.

Syntax `show openflow switch logical-switch-name flows`

Parameters `logical-switch-name`—Enter the logical switch instance name to view flow information.

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show openflow switch of-switch-1 flows
Logical switch name: of-switch-1
Total flows: 1
  Flow: 0
   Table ID: 0, Table: Ingress ACL TCAM table
   Flow ID: 0
   Priority: 32768, Cookie: 0
   Hard Timeout: 0, Idle Timeout: 0
   Packets: 0, Bytes: 0
   Match Parameters:
    In Port: ethernet1/1/1
    EType: 0x800
    SMAC: 00:0b:c4:a8:22:b0/ff:ff:ff:ff:ff:ff
    DMAC: 00:0b:c4:a8:22:b1/ff:ff:ff:ff:ff:ff
    VLAN id: 2/4095
    VLAN PCP: 1
    IP DSCP: 4
    IP ECN: 1
    IP Proto: 1
    Src Ip: 10.0.0.1/255.255.255.255
    Dst Ip: 20.0.0.1/255.255.255.255
    ICMPv4 Type: 1
    ICMPv4 Code: 10
    L4 Src Port: *
    L4 Dst Port: *
  Apply-Actions:  Output= ethernet1/1/2, ethernet1/1/3:1
  Write-Actions:  Drop
```

Supported Releases 10.4.1 or later

show openflow ports

Displays the OpenFlow ports for a specific logical switch instance.

Syntax `show openflow switch logical-switch-name ports`

Parameters `logical-switch-name`—Enter the name of the logical switch instance to view port information.

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show openflow switch of-switch-1 ports
Logical switch name: of-switch-1
```

```
Interface Name of-port ID  Config-State  Link-State  SPEED  DUPLEX  AUTONEG
TYPE
ethernet1/1/1     1        PORT_UP(CLI)  LINK_UP     40GB   FD      YES
COPPER
ethernet1/1/2     5        PORT_UP(CLI)  LINK_UP     40GB   FD      YES
COPPER
ethernet1/1/3:1   9        PORT_UP(CLI)  LINK_UP     10GB   FD      NO
FIBER
ethernet1/1/3:2   10       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
FIBER
ethernet1/1/3:3   11       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
FIBER
ethernet1/1/3:4   12       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
FIBER
ethernet1/1/4     13       PORT_UP(CLI)  LINK_UP     40GB   FD      YES
COPPER
ethernet1/1/5:1   17       PORT_UP(CLI)  LINK_UP     10GB   FD      NO
FIBER
ethernet1/1/5:2   18       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
FIBER
ethernet1/1/5:3   19       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
FIBER
ethernet1/1/5:4   20       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
FIBER
ethernet1/1/6     21       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/7     25       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/8     29       PORT_UP(CLI)  LINK_DOWN   0MB    FD      YES
COPPER
ethernet1/1/9     33       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/10    37       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/11    41       PORT_UP(CLI)  LINK_UP     40GB   FD      YES
COPPER
ethernet1/1/12    45       PORT_UP(CLI)  LINK_UP     40GB   FD      YES
COPPER
ethernet1/1/13    49       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/14    53       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/15    57       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/16    61       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/17    65       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/18    69       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/19    73       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/20    77       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/21    81       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/22    85       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/23    89       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/24    93       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/25    97       PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
COPPER
ethernet1/1/26    101      PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
COPPER
ethernet1/1/27    105      PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/28    109      PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
```

```
NONE
ethernet1/1/29    113      PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/30    117      PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/31    121      PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
ethernet1/1/32    125      PORT_UP(CLI)  LINK_DOWN   0MB    FD      NO
NONE
```

Supported Releases 10.4.1 or later

show openflow switch

Displays OpenFlow parameters for the switch instance.

Syntax `show openflow switch`

Parameters None

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show openflow switch
Logical switch name: of-switch-1
   Internal switch instance ID: 0
   Config state: true
   Signal Version: negotiate
   Data plane: secure
   Max backoff (sec): 8
   Probe Interval (sec): 5
   DPID: 90:b1:1c:f4:a5:23
   Switch Name : of-switch-1
   Number of buffers: 0
   Number of tables: 1
    Table ID: 0
    Table name: Ingress ACL TCAM table
    Max entries: 1000
    Active entries: 0
    Lookup count: 0
    Matched count: 0
   Controllers:
    10.16.208.150:6633, Protocol: none,
    packet-in Rate limit (packet per second): 0
    packet-in Burst limit: 0
```

Supported Releases 10.4.1 or later

show openflow switch controllers

Displays OpenFlow controllers for a specific logical switch instance.

Syntax `show openflow switch` *`logical-switch-name`* `controllers`

Parameters `logical-switch-name`—Enter the name of the logical switch instance to query.

Default None

Command Mode	EXEC
Usage Information	None
Example	

```
OS10# show openflow switch of-switch-1 controllers
Logical switch name: of-switch-1
Total Controllers: 1
  Controller: 1
    Target: 10.16.208.150:6633
    Protocol: TCP
    Connected: NO
    Role: Equal
    Last_error: Network is unreachable
    State: BACKOFF
    sec_since_disconnect: 0
```

Supported Releases	10.4.1 or later

switch

Creates a logical switch instance or modifies an existing logical switch instance.

Syntax	`switch` *`logical-switch-name`*
Parameters	`logical-switch-name`—Enter the name of the logical switch instance that you want to create or modify, a maximum of 15 characters. OS10 supports only one instance of the logical switch.
Default	None
Command Mode	OPENFLOW CONFIGURATION
Usage Information	You must configure a controller for the logical switch instance. The logical switch instance is disabled by default. To establish a connection with the controller, enable the logical switch instance using the `no shutdown` command. All physical and logical interfaces in the switch are assigned to the configured logical switch. The `no` form of this command removes the logical switch instance.

NOTE: OS10 supports only one instance of the logical switch. If you attempt to create a second logical switch instance, the following message appears:

```
% Warning: Only one Switch instance is supported
```

Example	

```
OS10# config terminal
OS10 (config)# openflow
OS10 (config-openflow)# switch of-switch-1
OS10 (config-openflow-switch)# no shutdown
```

Supported Releases	10.4.1 or later

OpenFlow-only mode commands

When you configure the switch to OpenFlow-only mode, only the following commands are available; all other commands are disabled.

NOTE:

- The `ntp` subcommand under the `interface` command is not applicable when the switch is in OpenFlow mode.
- The `ip` and `ipv6` subcommands under the `interface` command are applicable only when you configure the interface as the management port using the `in-band-mgmt` command.
- The `ip` and `ipv6` commands must be used only in In-Band mode (using the `in-band-mgmt` command).

Table 58. Modes and CLI commands

Mode	Available CLI commands

CONFIGURATION	`aaa`
	`alias`
	`banner`
	`class-map`
	`clock`
	`control-plane`
	`crypto`
	`end`
	`eula-consent`
	`exec-timeout`
	`exit`
	`feature`
	`help`
	`host-description`
	`hostname`
	`interface`
	`ip`
	· `ip access-list`
	· `ip route`
	· `ip ssh`
	· `ip telnet`
	`ipv6`
	· `ip access-list`
	`line`
	`logging`
	`login`
	`management`
	`no`
	`ntp`
	`openflow`
	`password-attributes`
	`policy-map`

Mode	Available CLI commands

	`radius-server`
	`rest`
	`scale-profile`
	`support-assist`
	`system`
	`tacacs-server`
	`trust`
	`username`
	`userrole`

EXEC	All commands The following debug commands are not available: • `debug iscsi` • `debug radius` • `debug tacacs+`
LAG INTERFACE CONFIGURATION	LAG is not supported.
LOOPBACK INTERFACE CONFIGURATION	Loopback interface is not supported.

INTERFACE CONFIGURATION	`description`
	`end`
	`exit`
	`ip`
	`mtu`
	`negotiation`
	`ntp`
	`show`
	`shutdown`

VLAN INTERFACE CONFIGURATION	VLAN is not supported.

11

Access Control Lists

OS10 uses two types of access policies — hardware-based ACLs and software-based route-maps. Use an ACL to filter traffic and drop or forward matching packets. To redistribute routes that match configured criteria, use a route-map.

ACLs

ACLs are a filter containing criterion to match; for example, examine internet protocol (IP), transmission control protocol (TCP), or user datagram protocol (UDP) packets, and an action to take such as forwarding or dropping packets at the NPU. ACLs permit or deny traffic based on MAC and/or IP addresses. The number of ACL entries is hardware-dependent.

ACLs have only two actions — forward or drop. Route-maps not only permit or block redistributed routes but also modify information associated with the route when it is redistributed into another protocol. When a packet matches a filter, the device drops or forwards the packet based on the filter's specified action. If the packet does not match any of the filters in the ACL, the packet drops, an implicit deny. ACL rules do not consume hardware resources until you apply the ACL to an interface.

ACLs process in sequence. If a packet does not match the criterion in the first filter, the second filter applies. If you configure multiple hardware-based ACLs, filter rules apply on the packet content based on the priority numeric processing unit (NPU) rule.

Route maps

Route-maps are software-based protocol filtering redistributing routes from one protocol to another and used in decision criterion in route advertisements. A route-map defines which of the routes from the specified routing protocol redistributes into the target routing process, see Route-maps.

Route-maps which have more than one match criterion, two or more matches within the same route-map sequence, have different match commands. Matching a packet against this criterion is an AND operation. If no match is found in a route-map sequence, the process moves to the next route-map sequence until a match is found, or until there are no more sequences. When a match is found, the packet forwards and no additional route-map sequences process. If you include a continue clause in the route-map sequence, the next route-map sequence also processes after a match is found.

The S5148F-ON switch has the following limitations:

- ACL counter does not support byte count.
- ACL rule does not look up the next header for IPv6 packets.
- Layer 2 (L2) Egress ACL does not work for unknown unicast traffic.
- L2 User ACL has a higher priority than the Layer 3 (L3) User ACL.
- You cannot modify or extend the hardware table for each ACL type.
- In IPv6 packets, only the protocol number of first header is matched.
- The egress Deny ACL entry does not block soft-forwarded packets and CPU-originated ICMP packets.

IP ACLs

An ACL filters packets based on the:

- IP protocol number
- Source and destination IP address
- Source and destination TCP port number

- Source and destination UDP port number

For ACL, TCP, and UDP filters, match criteria on specific TCP or UDP ports. For ACL TCP filters, you can also match criteria on established TCP sessions.

When creating an ACL, the sequence of the filters is important. You can assign sequence numbers to the filters as you enter them or OS10 can assign numbers in the order you create the filters. The sequence numbers display in the `show running-configuration` and `show ip access-lists [in | out]` command output.

Ingress and egress hot-lock ACLs allow you to append or delete new rules into an existing ACL without disrupting traffic flow. Existing entries in the content-addressable memory (CAM) shuffle to accommodate the new entries. Hot-lock ACLs are enabled by default and support ACLs on all platforms.

NOTE: Hot-lock ACLs support ingress ACLs only.

MAC ACLs

MAC ACLs filter traffic on the header of a packet. This traffic filtering is based on:

Source MAC packet address	MAC address range—address mask in 3x4 dotted hexadecimal notation, and *any* to denote that the rule matches all source addresses.
Destination MAC packet address	MAC address range—address-mask in 3x4 dotted hexadecimal notation, and *any* to denote that the rule matches all destination addresses.
Packet protocol	Set by its `EtherType` field contents and assigned protocol number for all protocols.
VLAN ID	Set in the packet header
Class of service	Present in the packet header

IPv4/IPv6 and MAC ACLs apply separately for inbound and outbound packets. You can assign an interface to multiple ACLs, with a limit of one ACL per packet direction per ACL type.

Control-plane ACLs

OS10 offers control-plane ACLs to selectively restrict packets that are destined to the CPU port, thereby providing increased security. Control-plane ACLs offer:

- An option to protect the CPU from denial of service (DoS) attacks.
- Fine-grained control to allow or block traffic going to the CPU.

Control-plane ACLs apply on the front-panel and management ports. Control-plane ACLs are one of the following types:

- IP ACL
- IPv6 ACL
- MAC ACL

NOTE: MAC ACL is applied only on packets that enter through the front-panel ports.

There is no implicit deny rule. If none of the configured conditions match, the default behavior is to permit. If you need to deny traffic that does not match any of the configured conditions, explicitly configure a deny statement.

The control-plane ACL is mutually exclusive with VTY ACL, the management ACL. VTY ACL provides secure access for session connection protocols, such as SSH or TELNET; however, control-plane ACLs permit or deny any TCP or UDP, including SSH and TELNET sessions, from specific hosts and networks, and also filters both IPv4 and IPv6 traffic.

Configure control-plane ACL

To configure control-plane ACLs, use the existing ACL template and create the appropriate rules to permit or deny traffic as needed, similar to creating an access list for VTY ACLs. However, when you apply this control-plane ACL, you must apply it in CONTROL-PLANE mode instead of VTY mode. For example:

```
OS10# configure terminal
OS10(config)# control-plane
OS10(config-control-plane)# ip access-group acl_name in
```

where *acl_name* is the name of the control-plane ACL, a maximum of 140 characters.

NOTE: Apply control-plane ACLs on ingress traffic only.

Control-plane ACL qualifiers

This section lists the supported control-plane ACL rule qualifiers.

NOTE: OS10 supports only the qualifiers listed below. Ensure that you use only these qualifiers in ACL rules.

- IPv4 qualifiers:
  - `DST_IP`—Destination IP address
  - `SRC_IP`—Source IP address
  - `IP_TYPE`—IP type
  - `IP_PROTOCOL`—Protocols such as TCP, UDP, and so on
  - `L4_DST_PORT`—Destination port number
- IPv6 qualifiers:
  - `DST_IPv6`—Destination address
  - `SRC_IPv6`—Source address
  - `IP_TYPE`—IP Type; for example, IPv4 or IPv6
  - `IP_PROTOCOL`—TCP, UDP, and so on
  - `L4_DST_PORT`—Destination port
- MAC qualifiers:
  - `OUT_PORT`—Egress CPU port
  - `SRC_MAC`—Source MAC address
  - `DST_MAC`—Destination MAC address
  - `ETHER_TYPE`—Ethertype
  - `OUTER_VLAN_ID`—VLAN ID
  - `IP_TYPE`—IP type
  - `OUTER_VLAN_PRI`—DOT1P value

IP fragment handling

OS10 supports a configurable option to explicitly deny IP-fragmented packets, particularly for the second and subsequent packets. This option extends the existing ACL command syntax with the `fragments` keyword for all L3 rules:

- Second and subsequent fragments are allowed because you cannot apply a L3 rule to these fragments. If the packet is denied eventually, the first fragment must be denied and the packet as a whole cannot be reassembled.
- The system applies implicit permit for the second and subsequent fragment before the *implicit* deny.
- If you configure an *explicit* deny, the second and subsequent fragments do not hit the implicit permit rule for fragments.

IP fragments ACL

When a packet exceeds the maximum packet size, the packet is fragmented into a number of smaller packets that contain portions of the contents of the original packet. This packet flow begins with an initial packet that contains all of the L3 and Layer 4 (L4) header information contained in the original packet, and is followed by a number of packets that contain only the L3 header information.

This packet flow contains all of the information from the original packet distributed through packets that are small enough to avoid the maximum packet size limit. This provides a particular problem for ACL processing.

If the ACL filters based on L4 information, the non-initial packets within the fragmented packet flow will not match the L4 information, even if the original packet would have matched the filter. Because of this filtering, packets are not processed by the ACL.

The examples show denying second and subsequent fragments, and permitting all packets on an interface. These ACLs deny all second and subsequent fragments with destination IP 10.1.1.1, but permit the first fragment and non-fragmented packets with destination IP 10.1.1.1. The second example shows ACLs which permits all packets — both fragmented and non-fragmented — with destination IP 10.1.1.1.

Deny second and subsequent fragments

```
OS10(config)# ip access-list ABC
OS10(conf-ipv4-acl)# deny ip any 10.1.1.1/32 fragments
OS10(conf-ipv4-acl)# permit ip any 10.1.1.1/32
```

Permit all packets on interface

```
OS10(config)# ip access-list ABC
OS10(conf-ipv4-acl)# permit ip any 10.1.1.1/32
OS10(conf-ipv4-acl)# deny ip any 10.1.1.1/32 fragments
```

L3 ACL rules

Use ACL commands for L3 packet filtering. TCP packets from host 10.1.1.1 with the TCP destination port equal to 24 are permitted, and all others are denied.

TCP packets that are first fragments or non-fragmented from host 10.1.1.1 with the TCP destination port equal to 24 are permitted, and all TCP non-first fragments from host 10.1.1.1 are permitted. All other IP packets that are non-first fragments are denied.

Permit ACL with L3 information only

If a packet's L3 information matches the information in the ACL, the packet's fragment offset (FO) is checked:

- If a packet's FO > 0, the packet is permitted
- If a packet's FO = 0, the next ACL entry processes

Deny ACL with L3 information only

If a packet's L3 information does not match the L3 information in the ACL, the packet's FO is checked:

- If a packet's FO > 0, the packet is denied
- If a packet's FO = 0, the next ACL line processes

Permit all packets from host

```
OS10(config)# ip access-list ABC
OS10(conf-ipv4-acl)# permit tcp host 10.1.1.1 any eq 24
OS10(conf-ipv4-acl)# deny ip any any fragment
```

Permit only first fragments and non-fragmented packets from host

```
OS10(config)# ip access-list ABC
OS10(conf-ipv4-acl)# permit tcp host 10.1.1.1 any eq 24
OS10(conf-ipv4-acl)# permit tcp host 10.1.1.1 any fragment
OS10(conf-ipv4-acl)# deny ip any any fragment
```

To log all packets denied and to override the implicit deny rule and the implicit permit rule for TCP/ UDP fragments, use a similar configuration. When an ACL filters packets, it looks at the FO to determine whether it is a fragment:

- FO = 0 means it is either the first fragment or the packet is a non-fragment
- FO > 0 means it is the fragments of the original packet

Assign sequence number to filter

IP ACLs filter on source and destination IP addresses, IP host addresses, TCP addresses, TCP host addresses, UDP addresses, and UDP host addresses. Traffic passes through the filter by filter sequence. Configure the IP ACL by first entering IP ACCESS-LIST mode and then assigning a sequence number to the filter.

User-provided sequence number

- Enter IP ACCESS LIST mode by creating an IP ACL in CONFIGURATION mode.

  ```
  ip access-list access-list-name
  ```

- Configure a drop or forward filter in IPV4-ACL mode.

  ```
  seq sequence-number {deny | permit | remark} {ip-protocol-number | icmp | ip | protocol | tcp
  | udp} {source prefix | source mask | any | host} {destination mask | any | host ip-address}
  [count [byte]] [fragments]
  ```

Auto-generated sequence number

If you are creating an ACL with only one or two filters, you can let the system assign a sequence number based on the order you configure the filters. The system assigns sequence numbers to filters using multiples of ten values.

- Configure a deny or permit filter to examine IP packets in IPV4-ACL mode.

  ```
  {deny | permit} {source mask | any | host ip-address} [count [byte]] [fragments]
  ```

- Configure a deny or permit filter to examine TCP packets in IPV4-ACL mode.

  ```
  {deny | permit} tcp {source mask] | any | host ip-address}} [count [byte]] [fragments]
  ```

- Configure a deny or permit filter to examine UDP packets in IPV4-ACL mode.

  ```
  {deny | permit} udp {source mask | any | host ip-address}} [count [byte]] [fragments]
  ```

Assign sequence number to filter

```
OS10(config)# ip access-list acl1
OS10(conf-ipv4-acl)# seq 5 deny tcp any any capture session 1 count
```

View ACLs and packets processed through ACL

```
OS10# show ip access-lists in
Ingress IP access-list acl1
 Active on interfaces :
  ethernet1/1/5
 seq 5 permit ip any any count (10000 packets)
```

Delete ACL rule

Before release 10.4.2, deleting ACL rules required a sequence number.

After release 10.4.2 or later, you can also delete ACL rules using the `no` form of the CLI command without using a sequence number.

While deleting ACL rules, the following conditions apply:

- Enter the exact `no` form of the CLI command. Each ACL rule is an independent entity. For example, the rule, `deny ip any any` is different from `deny ip any any count`.

  For example, if you configured the following rules:

  ```
  deny ip 1.1.1.1/24 2.2.2.2/24
  deny ip any any
  ```

  Using the `no deny ip any any` command deletes only the `deny ip any any` rule.

  To delete the `deny ip 1.1.1.1/24 2.2.2.2/24` rule, you must explicitly use the `no deny ip 1.1.1.1/24 2.2.2.2/24` command.

  NOTE: Wildcard option is not supported.
- You can no longer configure the same ACL rule multiple times using different sequence numbers. This option prevents duplicate rules from being entered in the system and taking up memory space.
- When you upgrade from a previous release to release 10.4.2 or later, the upgrade procedure removes all duplicate ACL rules and only one instance of an ACL rule remains in the system.

L2 and L3 ACLs

Configure both L2 and L3 ACLs on an interface in L2 mode. Rules apply if you use both L2 and L3 ACLs on an interface.

- L3 ACL filters packets and then the L2 ACL filters packets
- Egress L3 ACL filters packets

Rules apply in order:

- Ingress L3 ACL
- Ingress L2 ACL
- Egress L3 ACL
- Egress L2 ACL

NOTE: In ingress ACLs, L2 has a higher priority than L3 and in egress ACLs, L3 has a higher priority than L2.

Table 59. L2 and L3 targeted traffic

L2 ACL / L3 ACL	Targeted traffic
Deny / Deny	L3 ACL denies
Deny / Permit	L3 ACL permits
Permit / Deny	L3 ACL denies
Permit / Permit	L3 ACL permits

Assign and apply ACL filters

To filter an Ethernet interface, a port-channel interface, or a VLAN, assign an IP ACL filter to a physical interface. The IP ACL applies to all traffic entering a physical or port-channel interface. The traffic either forwards or drops depending on the criteria and actions you configure in the ACL filter.

To change the ACL filter functionality, apply the same ACL filters to different interfaces. For example, take ACL "ABCD" and apply it using the `in` keyword and it becomes an ingress ACL. If you apply the same ACL filter using the `out` keyword, it becomes an egress ACL.

You can apply an IP ACL filter to a physical or port-channel interface. The number of ACL filters allowed is hardware-dependent.

1 Enter the interface information in CONFIGURATION mode.

```
interface ethernet node/slot/port
```

2 Configure an IP address for the interface, placing it in L3 mode in INTERFACE mode.

```
ip address ip-address
```

3 Apply an IP ACL filter to traffic entering or exiting an interface in INTERFACE mode.

```
ip access-group access-list-name {in | out}
```

Configure IP ACL

```
OS10(config)# interface ethernet 1/1/28
OS10(conf-if-eth1/1/28)# ip address 10.1.2.0/24
OS10(conf-if-eth1/1/28)# ip access-group abcd in
```

View ACL filters applied to interface

```
OS10# show ip access-lists in
Ingress IP access-list acl1
Active on interfaces :
 ethernet1/1/28
seq 10 permit ip host 10.1.1.1 host 100.1.1.1 count (0 packets)
seq 20 deny ip host 20.1.1.1 host 200.1.1.1 count (0 packets)
seq 30 permit ip 10.1.2.0/24 100.1.2.0/24 count (0 packets)
seq 40 deny ip 20.1.2.0/24 200.1.2.0/24 count (0 packets)
seq 50 permit ip 10.0.3.0 255.0.255.0 any count (0 packets)
seq 60 deny ip 20.0.3.0 255.0.255.0 any count (0 packets)
seq 70 permit tcp any eq 1000 100.1.4.0/24 eq 1001 count (0 packets)
seq 80 deny tcp any eq 2100 200.1.4.0/24 eq 2200 count (0 packets)
seq 90 permit udp 10.1.5.0/28 eq 10000 any eq 10100 count (0 packets)
seq 100 deny tcp host 20.1.5.1 any rst psh count (0 packets)
seq 110 permit tcp any any fin syn rst psh ack urg count (0 packets)
seq 120 deny icmp 20.1.6.0/24 any fragment count (0 packets)
seq 130 permit 150 any any dscp 63 count (0 packets)
```

To view the number of packets matching the ACL, use the `count` option when creating ACL entries.

- Create an ACL that uses rules with the count option, see Assign sequence number to filter.

- Apply the ACL as an inbound or outbound ACL on an interface in CONFIGURATION mode, and view the number of packets matching the ACL.

```
show ip access-list {in | out}
```

Ingress ACL filters

To create an ingress ACL filter, use the `ip access-group` command in EXEC mode. To configure ingress, use the `in` keyword. Apply rules to the ACL with the `ip access-list` *`acl-name`* command. To view the access-list, use the `show access-lists` command.

1 Apply an ingress access-list on the interface in INTERFACE mode.

```
ip access-group access-group-name in
```

2 Return to CONFIGURATION mode.

```
exit
```

3 Create the access-list in CONFIGURATION mode.

```
ip access-list access-list-name
```

4 Create the rules for the access-list in ACCESS-LIST mode.

```
permit ip host ip-address host ip-address count
```

Apply ACL rules to access-group and view access-list

```
OS10(config)# interface ethernet 1/1/28
OS10(conf-if-eth1/1/28)# ip access-group abcd in
OS10(conf-if-eth1/1/28)# exit
OS10(config)# ip access-list acl1
OS10(conf-ipv4-acl)# permit ip host 10.1.1.1 host 100.1.1.1 count
```

Egress ACL filters

Egress ACL filters affect the traffic *leaving* the network. Configuring egress ACL filters onto physical interfaces protects the system infrastructure from a malicious and intentional attack by explicitly allowing only authorized traffic. These system-wide ACL filters eliminate the need to apply ACL filters onto each interface.

You can use an egress ACL filter to restrict egress traffic. For example, when you isolate denial of service (DoS) attack traffic to a specific interface, and apply an egress ACL filter to block the DoS flow from exiting the network, you protect downstream devices.

1 Apply an egress access-list on the interface in INTERFACE mode.

```
ip access-group access-group-name out
```

2 Return to CONFIGURATION mode.

```
exit
```

3 Create the access-list in CONFIGURATION mode.

```
ip access-list access-list-name
```

4 Create the rules for the access-list in ACCESS-LIST mode.

```
seq 10 deny ip any any count fragment
```

Apply rules to ACL filter

```
OS10(config)# interface ethernet 1/1/29
OS10(conf-if-eth1/1/29)# ip access-group egress out
OS10(conf-if-eth1/1/29)# exit
OS10(config)# ip access-list egress
OS10(conf-ipv4-acl)# seq 10 deny ip any any count fragment
```

View IP ACL filter configuration

```
OS10# show ip access-lists out
Egress IP access-list abcd
 Active on interfaces :
```

```
 ethernet1/1/29
seq 10 deny ip any any fragment count (100 packets)
```

Clear access-list counters

Clear IPv4, IPv6, or MAC access-list counters for a specific access-list or all lists. The counter counts the number of packets that match each permit or deny statement in an access-list. To get a more recent count of packets matching an access-list, clear the counters to start at zero. If you do not configure an access-list name, all IP access-list counters clear.

To view access-list information, use the `show access-lists` command.

- Clear IPv4 access-list counters in EXEC mode.
  `clear ip access-list counters access-list-name`
- Clear IPv6 access-list counters in EXEC mode.
  `clear ipv6 access-list counters access-list-name`
- Clear MAC access-list counters in EXEC mode.
  `clear mac access-list counters access-list-name`

IP prefix-lists

IP prefix-lists control the routing policy. An IP prefix-list is a series of sequential filters that contain a matching criterion and an permit or deny action to process routes. The filters process in sequence so that if a route prefix does not match the criterion in the first filter, the second filter applies, and so on.

A route prefix is an IP address pattern that matches on bits within the IP address. The format of a route prefix is `A.B.C.D/x`, where `A.B.C.D` is a dotted-decimal address and `/x` is the number of bits that match the dotted decimal address.

When the route prefix matches a filter, the system drops or forwards the packet based on the filter's designated action. If the route prefix does not match any of the filters in the prefix-list, the route drops, an implicit deny.

For example, in `112.24.0.0/16`, the first 16 bits of the address `112.24.0.0` match all addresses between `112.24.0.0` to `112.24.255.255`. Use permit or deny filters for specific routes with the `le` (less or equal) and `ge` (greater or equal) parameters, where `x.x.x.x/x` represents a route prefix:

- To deny only `/8` prefixes, enter `deny x.x.x.x/x ge 8 le 8`
- To permit routes with the mask greater than `/8` but less than `/12`, enter `permit x.x.x.x/x ge 8 le 12`
- To deny routes with a mask less than `/24`, enter `deny x.x.x.x/x le 24`
- To permit routes with a mask greater than `/20`, enter `permit x.x.x.x/x ge 20`

The following rules apply to prefix-lists:

- A prefix-list without permit or deny filters allows all routes
- An *implicit deny* is assumed — the route drops for all route prefixes that do not match a permit or deny filter
- After a route matches a filter, the filter's action applies and no additional filters apply to the route

Use prefix-lists in processing routes for routing protocols such as open shortest path first (OSPF), route table manager (RTM), and border gateway protocol (BGP).

To configure a prefix-list, use commands in PREFIX-LIST and ROUTER-BGP modes. Create the prefix-list in PREFIX-LIST mode and assign that list to commands in ROUTER-BGP modes.

Route-maps

Route-maps are a series of commands that contain a matching criterion and action. They change the packets meeting the matching criterion. ACLs and prefix-lists can only drop or forward the packet or traffic while route-maps process routes for route redistribution. For example, use a route-map to filter only specific routes and to add a metric.

- Route-maps also have an *implicit deny*. Unlike ACLs and prefix-lists where the packet or traffic drops, if a route does not match the route-map conditions, the route does not redistribute.
- Route-maps process routes for route redistribution. For example, to add a metric, a route-map can *filter* only specific routes. If the route does not match the conditions, the route-map decides where the packet or traffic drops. The route does not redistribute if it does not match.
- Route-maps use commands to decide what to do with traffic. To remove the match criteria in a route-map, use the `no match` command.
- In a BGP route-map, if you repeat the same match statements; for example, a match metric, with different values in the same sequence number, only the last match and set values are taken into account.

Configure match metric

```
OS10(config)# route-map hello
OS10(conf-route-map)# match metric 20
```

View route-map

```
OS10(conf-route-map)# do show route-map
route-map hello, permit, sequence 10
  Match clauses:
    metric 20
```

Change match

```
OS10(conf-route-map)# match metric 30
```

View updated route-map

```
OS10(conf-route-map)# do show route-map
route-map hello, permit, sequence 10
  Match clauses:
    metric 30
```

To filter the routes for redistribution, combine route-maps and IP prefix lists. If the route or packet matches the configured criteria, OS10 processes the route based on the `permit` or `deny` configuration of the prefix list.

When a route-map and a prefix list combine:

- For a route map with the `permit` action:
  - If a route matches a prefix-list set to `deny`, the route is denied
  - If a route matches a prefix-list set to `permit`, the route is permitted and any set of actions apply
- For a route map with the `deny` action:
  - If a route matches a prefix-list set to `deny`, the route is denied
  - If a route matches a prefix-list set to `permit`, the route is permitted and any set of actions apply

View both IP prefix-list and route-map configuration

```
OS10(conf-router-bgp-neighbor-af)# do show ip prefix-list
ip prefix-list p1:
seq 1 deny 10.1.1.0/24
seq 10 permit 0.0.0.0/0 le 32
ip prefix-list p2:
seq 1 permit 10.1.1.0/24
seq 10 permit 0.0.0.0/0 le 32
```

View route-map configuration

```
OS10(conf-router-bgp-neighbor-af)# do show route-map
route-map test1, deny, sequence 10
Match clauses:
ip address prefix-list p1
Set clauses:
route-map test2, permit, sequence 10
Match clauses:
ip address prefix-list p1
Set clauses:
route-map test3, deny, sequence 10
Match clauses:
ip address prefix-list p2
Set clauses:
route-map test4, permit, sequence 10
Match clauses:
ip address prefix-list p2
Set clauses:
```

Match routes

Configure match criterion for a route-map. There is no limit to the number of `match` commands per route map, but keep the number of match filters in a route-map low. The `set` commands do not require a corresponding `match` command.

- Match routes with a specific metric value in ROUTE-MAP mode, from 0 to 4294967295.
  `match metric metric-value`
- Match routes with a specific tag in ROUTE-MAP mode, from 0 to 4294967295.
  `match tag tag-value`
- Match routes whose next hop is a specific interface in ROUTE-MAP mode.
  `match interface interface`
  - `ethernet` — Enter the Ethernet interface information.
  - `port-channel` — Enter the port-channel number.
  - `vlan` — Enter the VLAN ID number.

Check match routes

```
OS10(config)# route-map test permit 1
0S10(conf-route-map)# match tag 250000
OS10(conf-route-map)# set weight 100
```

Set conditions

There is no limit to the number of `set` commands per route map, but keep the number of set filters in a route-map low. The `set` commands do not require a corresponding `match` command.

- Enter the IP address in A.B.C.D format of the next-hop for a BGP route update in ROUTE-MAP mode.
  `set ip next-hop address`
- Enter an IPv6 address in A::B format of the next-hop for a BGP route update in ROUTE-MAP mode.
  `set ipv6 next-hop address`
- Enter the range value for the BGP route's LOCAL_PREF attribute in ROUTE-MAP mode, from 0 to 4294967295.
  `set local-preference range-value`
- Enter a metric value for redistributed routes in ROUTE-MAP mode, from 0 to 4294967295.
  `set metric {+ | - | metric-value}`
- Enter an OSPF type for redistributed routes in ROUTE-MAP mode.
  `set metric-type {type-1 | type-2 | external | internal}`

- Enter an ORIGIN attribute in ROUTE-MAP mode.
  ```
  set origin {egp | igp | incomplete}
  ```
- Enter a tag value for the redistributed routes in ROUTE-MAP mode, from 0 to 4294967295.
  ```
  set tag tag-value
  ```
- Enter a value as the route's weight in ROUTE-MAP mode, from 0 to 65535.
  ```
  set weight value
  ```

Check set conditions

```
OS10(config)# route-map ip permit 1
OS10(conf-route-map)# match metric 2567
```

Continue clause

Only BGP route-maps support the `continue` clause. When a match is found, `set` clauses run and the packet forwards — no route-map processing occurs. If you configure the `continue` clause without configuring a module, the next sequential module processes.

If you configure the `continue` command at the end of a module, the next module processes even after a match is found. The example shows a continue clause at the end of a route-map module — if a match is found in the route-map `test` module 10, module 30 processes.

Route-map continue clause

```
OS10(config)# route-map test permit 10
OS10(conf-route-map)# continue 30
```

ACL flow-based monitoring

Flow-based monitoring conserves bandwidth by selecting only the required flow to mirror instead of mirroring entire packets from an interface. This feature is available for L2 and L3 ingress traffic. Specify flow-based monitoring using ACL rules. Flow-based monitoring copies incoming packets that match the ACL rules applied on the ingress port and forwards, or mirrors them to another port. The source port is the monitored port (MD), and the destination port is the monitoring port (MG).

When a packet arrives at a monitored port, the packet validates against the configured ACL rules. If the packet matches an ACL rule, the system examines the corresponding flow processor and performs the action specified for that port. If the mirroring action is set in the flow processor entry, the port details are sent to the destination port.

Flow-based mirroring

Flow-based mirroring is a mirroring session in which traffic matches specified policies that mirrors to a destination port. Port-based mirroring maintains a database that contains all monitoring sessions, including port monitor sessions. The database has information regarding the sessions that are enabled or not enabled for flow-based monitoring. Flow-based mirroring is also known as policy-based mirroring.

To enable flow-based mirroring, use the `flow-based enable` command. Traffic with particular flows that traverse through the ingress interfaces are examined. Appropriate ACL rules apply in the ingress direction. By default, flow-based mirroring is not enabled.

To enable evaluation and replication of traffic traversing to the destination port, configure the monitor option using the `permit`, `deny`, or `seq` commands for ACLs assigned to the source or the monitored port (MD). Enter the keywords `capture session session-id` with the `seq`, `permit`, or `deny` command for the ACL rules to allow or drop IPv4, IPv6, ARP, UDP, EtherType, ICMP, and TCP packets.

IPV4-ACL mode

```
seq sequence-number {deny | permit} {source [mask] | any | host ip-address} [count [byte]]
[fragments] [threshold-in-msgs count] [capture session session-id]
```

If you configure the `flow-based enable` command and do not apply an ACL on the source port or the monitored port, both flow-based monitoring and port mirroring do not function. Flow-based monitoring is supported only for ingress traffic.

The `show monitor session` *`session-id`* command displays output that indicates if a particular session is enabled for flow-monitoring.

View flow-based monitoring

```
OS10# show monitor session 1
S.Id  Source        Destination    Dir  SrcIP  DstIP  DSCP  TTL  State Reason
-------------------------------------------------------------------------------
1    ethernet1/1/1  ethernet1/1/4  both  N/A   N/A    N/A  N/A  true   Is UP
```

Traffic matching ACL rule

```
OS10# show ip access-lists in
Ingress IP access-list testflow
 Active on interfaces :
  ethernet1/1/1
 seq 5 permit icmp any any capture session 1
 seq 10 permit ip 102.1.1.0/24 any capture session 1
 seq 15 deny udp any any capture session 2
 seq 20 deny tcp any any capture session 3
```

Enable flow-based monitoring

Flow-based monitoring conserves bandwidth by mirroring only specified traffic, rather than all traffic on an interface. It is available for L2 and L3 ingress and egress traffic. Configure traffic to monitor using ACL filters.

1 Create a monitor session in MONITOR-SESSION mode.

```
monitor session session-number type local
```

2 Enable flow-based monitoring for the mirroring session in MONITOR-SESSION mode.

```
flow-based enable
```

3 Define ACL rules that include the keywords `capture session` *`session-id`* in CONFIGURATION mode. The system only considers port monitoring traffic that matches rules with the keywords `capture session`.

```
ip access-list
```

4 Apply the ACL to the monitored port in INTERFACE mode.

```
ip access-group access-list
```

Enable flow-based monitoring

```
OS10(config)# monitor session 1 type local
OS10(conf-mon-local-1)# flow-based enable
OS10(config)# ip access-list testflow
OS10(conf-ipv4-acl)# seq 5 permit icmp any any capture session 1
OS10(conf-ipv4-acl)# seq 10 permit ip 102.1.1.0/24 any capture session 1
OS10(conf-ipv4-acl)# seq 15 deny udp any any capture session 2
OS10(conf-ipv4-acl)# seq 20 deny tcp any any capture session 3
OS10(conf-ipv4-acl)# exit
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# ip access-group testflow in
OS10(conf-if-eth1/1/1)# no shutdown
```

View access-list configuration

```
OS10# show ip access-lists in
Ingress IP access-list testflow
 Active on interfaces :
  ethernet1/1/1
 seq 5 permit icmp any any capture session 1
 seq 10 permit ip 102.1.1.0/24 any capture session 1
 seq 15 deny udp any any capture session 2
 seq 20 deny tcp any any capture session 3
```

View monitor sessions

```
OS10(conf-if-eth1/1/1)# show monitor session all
S.Id  Source       Destination    Dir  SrcIP  DstIP  DSCP TTL  State  Reason
-------------------------------------------------------------------------
1   ethernet1/1/1  ethernet1/1/4  both  N/A   N/A    N/A  N/A  true   Is UP
```

View ACL table utilization report

The `show acl-table-usage detail` command shows the ingress and egress ACL tables for the various features and their utilization.

The hardware pool area displays the ingress application groups (pools), the features mapped to each of these groups, and the amount of used and free space available in each of the pools. The amount of space required to store a single ACL rule in a pool depends on the keywidth of the TCAM slice.

The service pool displays the amount of used and free space for each of the features. The number of ACL rules configured for a feature is displayed in the configured rules column. The number of used rows depends on the number of ports the configured rules are applied on. Under Allocated pools, you can view the percentage of dedicated space reserved for a particular feature or the phrase Shared if you have not reserved space for each of the features individually, against the total number of pools allocated for the application group. In the example given below, the SYSTEM_FLOW feature has 15 percentage of space reserved in ingress app-group-1 with a pool count of 1, which is represented by 15:1.

```
OS10# show acl-table-usage detail
Ingress ACL utilization
Hardware Pools
-------------------------------------------------------------------------------------------------------
-------
Pool ID     App(s)                                                    Used rows    Free rows
Max rows
-------------------------------------------------------------------------------------------------------
-------
0          SYSTEM_FLOW                                                 49           975
1024
1          SYSTEM_FLOW                                                 49           975
1024
2          USER_IPV4_ACL                                               3            1021
1024
3          USER_L2_ACL                                                 2            1022
1024
4          USER_IPV6_ACL                                               2            510
512
5          USER_IPV6_ACL                                               2            510
512
6          FCOE                                                        55           457
512
7          FCOE                                                        55           457
512
8          ISCSI_SNOOPING                                              12           500
512
9          FREE                                                        0            512
512
10         PBR_V6                                                      1            511
512
11         PBR_V6                                                      1            511
512
-------------------------------------------------------------------------------------------------------
-------
Service Pools
-------------------------------------------------------------------------------------------------------
-------
App                 Allocated pools  App group   Configured rules    Used rows    Free rows
Max rows
-------------------------------------------------------------------------------------------------------
-------
USER_L2_ACL         Shared:1         G3          1                   2            1022
```

```
1024
USER_IPV4_ACL       Shared:1         G2          2                   3            1021
1024
USER_IPV6_ACL       Shared:2         G4          1                   2            510
512
PBR_V6              Shared:2         G10         1                   1            511
512
SYSTEM_FLOW         Shared:2         G0          49                  49           975
1024
ISCSI_SNOOPING      Shared:1         G8          12                  12           500
512
FCOE                Shared:2         G6          55                  55           457
512
-------------------------------------------------------------------------------------------------
-------
Egress ACL utilization
Hardware Pools
-------------------------------------------------------------------------------------------------
-------
Pool ID     App(s)                                                  Used rows    Free rows
Max rows
-------------------------------------------------------------------------------------------------
-------
0         USER_IPV4_EGRESS                                           2            254
256
1         USER_L2_ACL_EGRESS                                         2            254
256
2         USER_IPV6_EGRESS                                           2            254
256
3         USER_IPV6_EGRESS                                           2            254
256
-------------------------------------------------------------------------------------------------
-------
Service Pools
-------------------------------------------------------------------------------------------------
-------
App                 Allocated pools  App group   Configured rules    Used rows    Free rows
Max rows
-------------------------------------------------------------------------------------------------
-------
USER_L2_ACL_EGRESS  Shared:1         G1          1                   2            254
256
USER_IPV4_EGRESS    Shared:1         G0          1                   2            254
256
USER_IPV6_EGRESS    Shared:2         G2          1                   2            254
256
```

Known behavior

- On the S4200-ON platform, the `show acl-table-usage detail` command output lists several hardware pools as available (FREE), but you will see an "ACL CAM table full" warning log when the system creates a new service pool. The system will not be able to create any new service pools. The existing groups, however, can continue to grow up to the maximum available pool space.
- On the S4200-ON platform, the `show acl-table usage detail` command output lists all the available hardware pools under Ingress ACL utilization table and none under the Egress ACL utilization table. The system allocates pool space for Egress ACL table only when you configure Egress ACLs. You can run the `show acl-table-usage detail` command again to view pool space allocated under Egress ACL utilization table as well.
- On S52xx-ON, Z91xx-ON, Z92xx-ON platforms, the number of Configured Rules listed under Service Pools for each of the features is the number of ACLs multiplied by the number of ports on which they are applied. This number is cumulative. You can view the Used rows and Free rows that indicate the actual amount of space that is utilized and available in the hardware.

ACL logging

You can configure ACLs to filter traffic, drop or forward packets that match certain conditions. The ACL logging feature allows you to get additional information about packets that match an access control list entry (ACE) applied on an interface in inbound direction.

OS10 creates a log message that includes additional information about the packet, when a matching packet hits a log-enabled ACL entry.

ACL logging helps to administer and manage traffic that traverses your network and is useful for network supervision and maintenance activities.

High volumes of network traffic can result in large volume of logs, which can negatively impact system performance and efficiency.

You can specify the threshold after which a log is created and the interval at which the logs must be created.

The threshold defines how often a log message is created after an initial packet match. The default is 10 messages. This value is configurable and the range is from 1 to 100 messages.

By default, the interval is set to 5 minutes and logs are created every 5 minutes. During this interval, the system continues to examine the packets against the configured ACL rule and permits or denies traffic, but logging is halted temporarily. This value is configurable and the range is from 1 to 10 minutes.

For example, if you have configured a threshold value of 20 and an interval of 10 minutes, after an initial packet match, the 20th packet that matches the ACL entry is logged. The system then waits for the interval period of 10 minutes to elapse, during which time no logging occurs. Once the interval period elapses, the 20th packet that matches the ACL entry is logged again.

NOTE: This feature is not supported on the S5148F-ON platform.

Important notes

The ACL logging feature is:

- Applicable only for IPv4 and IPv6 user ACLs. MAC ACLs are not logged.
- Applicable only for user ACLs applied on interfaces in the inbound direction. Even though ACL logging cannot be enabled for outbound ACLs, ACL configuration is applied.
- Not supported on the S5148F-ON platform.

If you have enabled ACL logging, downgrade from release 10.4.3.0 to a previous release fails and the corresponding ACL rules are not applied. Before you downgrade, be sure to disable ACL logging or delete the startup configuration.

Dell EMC recommends that you do not enable logging for control plane protocol identical user ACL entry.

DELL EMC recommends a max scale of 128 log-enabled ACL entries. If logging cannot be enabled on further ACL entries, a syslog error message appears to indicate logging cannot be enabled. However, the ACL entries are applied.

ACL commands

clear ip access-list counters

Clears ACL counters for a specific access-list.

Syntax `clear ip access-list counters [access-list-name]`

Parameters	*access-list-name* — (Optional) Enter the name of the IP access-list to clear counters. A maximum of 140 characters.
Default	Not configured
Command Mode	EXEC
Usage Information	If you do not enter an access-list name, all IPv6 access-list counters clear. The counter counts the number of packets that match each permit or deny statement in an access-list. To get a more recent count of packets matching an access list, clear the counters to start at zero. To view access-list information, use the `show access-lists` command.
Example	`OS10# clear ip access-list counters`
Supported Releases	10.2.0E or later

clear ipv6 access-list counters

Clears IPv6 access-list counters for a specific access-list.

Syntax	`clear ipv6 access-list counters [access-list-name]`
Parameters	*access-list-name* — (Optional) Enter the name of the IPv6 access-list to clear counters. A maximum of 140 characters.
Default	Not configured
Command Mode	EXEC
Usage Information	If you do not enter an access-list name, all IPv6 access-list counters clear. The counter counts the number of packets that match each permit or deny statement in an access list. To get a more recent count of packets matching an access list, clear the counters to start at zero. To view access-list information, use the `show access-lists` command.
Example	`OS10# clear ipv6 access-list counters`
Supported Releases	10.2.0E or later

clear mac access-list counters

Clears counters for a specific or all MAC access lists.

Syntax	`clear mac access-list counters [access-list-name]`
Parameters	*access-list-name* — (Optional) Enter the name of the MAC access list to clear counters. A maximum of 140 characters.
Default	Not configured
Command Mode	EXEC
Usage Information	If you do not enter an access-list name, all MAC access-list counters clear. The counter counts the number of packets that match each permit or deny statement in an access list. To get a more recent count of packets matching an access list, clear the counters to start at zero. To view access-list information, use the `show access-lists` command.
Example	`OS10# clear mac access-list counters`

Supported Releases 10.2.0E or later

deny

Configures a filter to drop packets with a specific IP address.

Syntax

```
deny [protocol-number | icmp | ip | tcp | udp] [A.B.C.D | A.B.C.D/x | any |
host ip-address] [A.B.C.D | A.B.C.D/x | any | host ip-address] [capture | dscp
value | fragment]
```

Parameters

- *protocol-number* — (Optional) Enter the protocol number identified in the IP header, from 0 to 255.
- icmp — (Optional) Enter the ICMP address to deny.
- ip — (Optional) Enter the IP address to deny.
- tcp — (Optional) Enter the TCP address to deny.
- udp — (Optional) Enter the UDP address to deny.
- A.B.C.D — Enter the IP address in dotted decimal format.
- A.B.C.D/x — Enter the number of bits to match to the dotted decimal address.
- any — (Optional) Enter the filter type to subject routes to.
  - capture — (Optional) Capture packets the filter processes.
  - dscp *value* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
  - fragment — (Optional) Use ACLs to control packet fragments.
- host *ip-address* — (Optional) Enter the keyword and the IP address to use a host address only.

Default Not configured

Command Mode IPV4-ACL

Usage Information The no version of this command removes the filter.

Example

```
OS10(config)# ip access-list testflow
OS10(conf-ipv4-acl)# deny udp any any
```

Supported Releases 10.2.0E or later

deny (IPv6)

Configures a filter to drop packets with a specific IPv6 address.

Syntax

```
deny [protocol-number | icmp | ipv6 | tcp | udp] [A::B | A::B/x | any | host
ipv6-address] [A::B | A::B/x | any | host ipv6-address] [capture | dscp value |
fragment]
```

Parameters

- *protocol-number* — (Optional) Enter the protocol number identified in the IP header, from 0 to 255.
- icmp — (Optional) Enter the ICMP address to deny.
- ipv6 — (Optional) Enter the IPv6 address to deny.
- tcp — (Optional) Enter the TCP address to deny.
- udp — (Optional) Enter the UDP address to deny.
- A::B — Enter the IPv6 address in dotted decimal format.
- A::B/x — Enter the number of bits to match to the IPv6 address.

- any — (Optional) Enter so that all routes are subject to the filter:
  - capture — (Optional) Capture packets the filter processes.
  - dscp *value* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
  - fragment — (Optional) Use ACLs to control packet fragments.
- host *ipv6-address* — (Optional) Enter the keyword and the IPv6 address to use a host address only.

Default Not configured

Command Mode IPV6-ACL

Usage Information The no version of this command removes the filter.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# deny ipv6 any any capture session 1
```

Supported Releases 10.2.0E or later

deny (MAC)

Configures a filter to drop packets with a specific MAC address.

Syntax

```
deny {nn:nn:nn:nn:nn:nn [00:00:00:00:00:00] | any} {nn:nn:nn:nn:nn:nn
[00:00:00:00:00:00] | any} [protocol-number | capture | cos | vlan]
```

Parameters

- *nn:nn:nn:nn:nn:nn* — Enter the MAC address of the network from or to which the packets are sent.
- *00:00:00:00:00:00* — (Optional) Enter which bits in the MAC address must match. If you do not enter a mask, a mask of 00:00:00:00:00:00 applies.
- any — (Optional) Set routes which are subject to the filter.
  - *protocol-number* — (Optional) MAC protocol number identified in the header, from 600 to ffff.
  - capture — (Optional) Capture packets the filter processes.
  - cos — (Optional) CoS value, from 0 to 7.
  - vlan — (Optional) VLAN number, from 1 to 4093.

Default Disabled

Command Mode MAC-ACL

Usage Information The no version of this command removes the filter.

Example

```
OS10(config)# mac access-list macacl
OS10(conf-mac-acl)# deny any any cos 7
OS10(conf-mac-acl)# deny any any vlan 2
```

Supported Releases 10.2.0E or later

deny icmp

Configures a filter to drop all or specific Internet Control Message Protocol (ICMP) messages.

Syntax

```
deny icmp [A.B.C.D | A.B.C.D/x | any | host ip-address] [[A.B.C.D | A.B.C.D/x |
any | host ip-address] [capture | dscp value | fragment]
```

Parameters

- `A.B.C.D` — Enter the IP address in hexadecimal format separated by colons.
- `A.B.C.D/x` — Enter the number of bits to match to the IP address.
- `any` — (Optional) Set all routes subject to the filter.
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp value` — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
- `host ip-address` — (Optional) Enter the IP address to use a host address only.

Default Not configured

Command Mode IPV4-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# ip access-list egress
OS10(conf-ipv4-acl)# deny icmp any any capture session 1
```

Supported Releases 10.2.0E or later

deny icmp (IPv6)

Configures a filter to drop all or specific ICMP messages.

Syntax

```
deny icmp [A::B | A::B/x | any | host ipv6-address] [A::B | A::B/x | any | host
ipv6-address] [capture | dscp value | fragment]
```

Parameters

- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits to match to the IPv6 address.
- `any` — (Optional) Set all routes which are subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp value` — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
- `host ipv6-address` — (Optional) Enter the IPv6 address to use a host address only.

Default Not configured

Command Mode IPV6-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# deny icmp any any capture session 1
```

Supported Releases 10.2.0E or later

deny ip

Configures a filter to drop all or specific packets from an IPv4 address.

Syntax

```
deny ip [A.B.C.D | A.B.C.D/x | any | host ip-address] [[A.B.C.D | A.B.C.D/x |
any | host ip-address] [capture | dscp value | fragment]
```

Parameters

- `A.B.C.D` — Enter the IPv4 address in dotted decimal format.
- `A.B.C.D/x` — Enter the number of bits to match to the dotted decimal address.
- `any` — (Optional) Set all routes which are subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp` *`value`* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
- `host` *`ip-address`* — (Optional) Enter the IPv4 address to use a host address only.

Default Not configured

Command Mode IPV4-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# ip access-list testflow
OS10(conf-ipv4-acl)# deny ip any any capture session 1 count
```

Supported Releases 10.2.0E or later

deny ipv6

Configures a filter to drop all or specific packets from an IPv6 address.

Syntax

```
deny ipv6 [A::B | A::B/x | any | host ipv6-address] [A::B | A:B/x | any | host
ipv6-address] [capture | dscp | fragment]
```

Parameters

- `A::B` — (Optional) Enter the source IPv6 address from which the packet was sent and the destination address.
- `A::B/x` — (Optional) Enter the source network mask in /prefix format (/x) and the destination mask.
- `any` — (Optional) Set all routes which are subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp` *`value`* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
- `host` *`ipv6-address`* — (Optional) Enter the IPv6 address to use a host address only.

Default Not configured

Command Mode IPV6-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# deny ipv6 any any capture session 1
```

Supported Releases 10.2.0E or later

deny tcp

Configures a filter that drops Transmission Control Protocol (TCP) packets meeting the filter criteria.

Syntax

```
deny tcp [A.B.C.D | A.B.C.D/x | any | host ip-address [operator]] [[A.B.C.D |
A.B.C.D/x | any | host ip-address [operator]] [ack | fin | psh | rst | syn |
urg] [capture | dscp value | fragment]
```

Parameters

- `A.B.C.D` — Enter the IPv4 address in A.B.C.D format.
- `A.B.C.D/x` — Enter the number of bits to match in A.B.C.D/x format.
- `any` — (Optional) Enter to subject all routes to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp value` — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
  - `ack` — (Optional) Set the bit as acknowledgement.
  - `fin` — (Optional) Set the bit as finish—no more data from sender.
  - `psh` — (Optional) Set the bit as push.
  - `rst` — (Optional) Set the bit as reset.
  - `syn` — (Optional) Set the bit as synchronize.
  - `urg` — (Optional) Set the bit set as urgent.
- *operator* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - `eq` — Equal to
  - `gt` — Greater than
  - `lt` — Lesser than
  - `neq` — Not equal to
  - `range` — Range of ports, including the specified port numbers.
- `host ip-address` — (Optional) Enter the keyword and the IPv4 address to use a host address only.

Default Not configured

Command Mode IPV4-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# ip access-list testflow
OS10(conf-ipv4-acl)# deny tcp any any capture session 1
```

Supported Releases 10.2.0E or later

deny tcp (IPv6)

Configures a filter that drops TCP IPv6 packets meeting the filter criteria.

Syntax

```
deny tcp [A::B | A::B/x | any | host ipv6-address [operator]] [A::B | A:B/x |
any | host ipv6-address [operator]] [ack | fin | psh | rst | syn | urg]
[capture | dscp value | fragment]
```

Parameters

- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits to match to the IPv6 address.
- `any` — (Optional) Set all routes which are subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp` *`value`* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
- *`operator`* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - `eq` — Equal to
  - `gt` — Greater than
  - `lt` — Lesser than
  - `neq` — Not equal to
  - `range` — Range of ports, including the specified port numbers.
- `host` *`ipv6-address`* — (Optional) Enter the IPv6 address to use a host address only.

Default Not configured

Command Mode IPV6-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# deny tcp any any capture session 1
```

Supported Releases 10.2.0E or later

deny udp

Configures a filter to drop User Datagram Protocol (UDP) packets meeting the filter criteria.

Syntax

```
deny udp [A.B.C.D | A.B.C.D/x | any | host ip-address [operator]] [A.B.C.D |
A.B.C.D/x | any | host ip-address [operator]] [ack | fin | psh | rst | syn |
urg] [capture | dscp value | fragment]
```

Parameters

- `A.B.C.D` — Enter the IPv4 address in dotted decimal format.
- `A.B.C.D/x` — Enter the number of bits to match to the dotted decimal address.
- `any` — (Optional) Set all routes which are subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp` *`value`* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
  - `ack` — (Optional) Set the bit as acknowledgement.
  - `fin` — (Optional) Set the bit as finish—no more data from sender.
  - `psh` — (Optional) Set the bit as push.
  - `rst` — (Optional) Set the bit as reset.
  - `syn` — (Optional) Set the bit as synchronize.
  - `urg` — (Optional) Set the bit set as urgent.
- *`operator`* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - `eq` — Equal to

  - `gt` — Greater than
  - `lt` — Lesser than
  - `neq` — Not equal to
  - `range` — Range of ports, including the specified port numbers.
- `host` *`ip-address`* — (Optional) Enter the IPv4 address to use a host address only.

Default Not configured

Command Mode IPV4-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# ip access-list testflow
OS10(conf-ipv4-acl)# deny udp any any capture session 1
```

Supported Releases 10.2.0E or later

deny udp (IPv6)

Configures a filter to drop UDP IPv6 packets that match filter criteria.

Syntax

```
deny udp [A::B | A::B/x | any | host ipv6-address [operator]] [A::B | A:B/x |
any | host ipv6-address [operator]] [ack | fin | psh | rst | syn | urg]
[capture | dscp value | fragment]
```

Parameters

- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits to match to the IPv6 address.
- `any` — (Optional) Enter for all routes to be subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp` *`value`* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
  - `ack` — (Optional) Set the bit as acknowledgement.
  - `fin` — (Optional) Set the bit as finish—no more data from sender.
  - `psh` — (Optional) Set the bit as push.
  - `rst` — (Optional) Set the bit as reset.
  - `syn` — (Optional) Set the bit as synchronize.
  - `urg` — (Optional) Set the bit set as urgent.
- *`operator`* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - `eq` — Equal to
  - `gt` — Greater than
  - `lt` — Lesser than
  - `neq` — Not equal to
  - `range` — Range of ports, including the specified port numbers.
- `host` *`ipv6-address`* — (Optional) Enter the keyword and the IPv6 address to use a host address only.

Default Not configured

Command Mode IPV6-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# deny udp any any capture session 1
```

Supported Releases 10.2.0E or later

description

Configures an ACL description.

Syntax `description text`

Parameters *text* — Enter the description text string. A maximum of 80 characters.

Default Disabled

Command Modes IPV4-ACL, IPV6-ACL, MAC-ACL

Usage Information The `no` version of this command deletes the ACL description.

Example

```
OS10(conf-ipv4-acl)# description ipacltest
```

Supported Releases 10.2.0E or later

ip access-group

Configures an IPv4 access group.

Syntax `ip access-group access-list-name {in | out}`

Parameters

- *access-list-name* — Enter the name of an IPv4 access list. A maximum of 140 characters.
- `in` — Apply the ACL to incoming traffic.
- `out` — Apply the ACL to outgoing traffic.

Default Not configured

Command Mode INTERFACE

CONTROL-PLANE

Usage Information Use this command in the CONTROL-PLANE mode to apply a control-plane ACL. Control-plane ACLs are only applied on the ingress traffic. By default, the control-plane ACL is applied to the front-panel ports as well as the management port.The `no` version of this command deletes the IPv4 ACL configuration.

Example

```
OS10(conf-if-eth1/1/8)# ip access-group testgroup in
```

Example (Control-plane ACL)

```
OS10# configure terminal
OS10(config)# control-plane
OS10(config-control-plane)# ip access-group aaa-cp-acl in
```

Supported Releases 10.2.0E or later; 10.4.1 or later (control-plane ACL)

ip access-list

Creates an IP access list to filter based on an IP address.

Syntax `ip access-list access-list-name`

Parameters *access-list-name* — Enter the name of an IPv4 access list. A maximum of 140 characters.

Default Not configured

Command Mode CONFIGURATION

Usage Information None

Example

```
OS10(config)# ip access-list acl1
```

Supported Releases 10.2.0E or later

ip as-path access-list

Create an AS-path ACL filter for BGP routes using a regular expression.

Syntax `ip as-path access-list name {deny | permit} regexp-string`

Parameters

- *name* — Enter an access list name.
- `deny | permit` — Reject or accept a matching route.
- *regexp-string* — Enter a regular expression string to match an AS-path route attribute.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information You can specify an access-list filter on inbound and outbound BGP routes. The ACL filter consists of regular expressions. If a regular expression matches an AS path attribute in a BGP route, the route is rejected or accepted. The AS path does not contain the local AS number. The `no` version of this command removes a single access list entry if you specify `deny` and a `regexp`. Otherwise, the entire access list is removed.

The `question mark (?)` character is not supported in the regular expressions. All other special characters are supported. When you are using `backslash(\)` or `double quotes (")` in the regular expression, precede these characters with `backslash(\)`. For example, enter `\\` or `\"`.

Example

```
OS10(config)# ip as-path access-list abc deny 123
```

Supported Release 10.3.0E or later

ip community-list standard deny

Creates a standard community list for BGP to deny access.

Syntax `ip community-list standard name deny {aa:nn | no-advertise | local-AS | no-export | internet}`

Parameters

- `name` — Enter the name of the standard community list used to identify one more deny groups of communities.
- *`aa:nn`* — Enter the community number in the format *aa:nn*, where *aa* is the number that identifies the autonomous system and *nn* is a number the identifies the community within the autonomous system.
- `no-advertise` — BGP does to not advertise this route to any internal or external peer.
- `local-AS` — BGP does not advertise this route to external peers.
- `no-export` — BGP does not advertise this route outside a BGP confederation boundary.
- `internet` — BGP does not advertise this route to an Internet community.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the community list.

Example

```
OS10(config)# ip community-list standard STD_LIST deny local-AS
```

Supported Release 10.3.0E or later

ip community–list standard permit

Creates a standard community list for BGP to permit access.

Syntax

```
ip community-list standard name permit {aa:nn | no-advertise | local-as | no-export | internet}
```

Parameters

- `name` — Enter the name of the standard community list used to identify one more deny groups of communities.
- *`aa:nn`* — Enter the community number in the format *aa:nn*, where *aa* is the number that identifies the autonomous system and *nn* is a number the identifies the community within the autonomous system.
- `no-advertise` — BGP does not advertise this route to any internal or external peer.
- `local-as` — BGP does not advertise this route to external peers.
- `no-export` — BGP does not advertise this route outside a BGP confederation boundary
- `internet` — BGP does not advertise this route to an Internet community.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the community list.

Example

```
OS10(config)# ip community-list standard STD_LIST permit local-AS
```

Supported Release 10.3.0E or later

ip extcommunity-list standard deny

Creates an extended community list for BGP to deny access.

Syntax

```
ip extcommunity-list standard name deny {4byteas-generic | rt | soo}
```

Parameters

- *name* — Enter the name of the community list used to identify one or more deny groups of extended communities.
- 4byteas-generic—Enter the generic extended community then the keyword transitive or non-transitive.
- *rt* — Enter the route target.
- *soo* — Enter the route origin or site-of-origin.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The no version of this command removes the extended community list.

Example

```
OS10(config)# ip extcommunity-list standard STD_LIST deny 4byteas-generic
transitive 1.65534:40
```

Supported Release 10.3.0E or later

ip extcommunity-list standard permit

Creates an extended community list for BGP to permit access.

Syntax `ip extcommunity-list standard name permit {4byteas-generic | rt | soo}`

Parameters

- *name* — Enter the name of the community list used to identify one or more permit groups of extended communities.
- *rt* — Enter the route target.
- *soo* — Enter the route origin or site-of-origin.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The no version of this command removes the extended community list.

Example

```
OS10(config)# ip extcommunity-list standard STD_LIST permit 4byteas-generic
transitive 1.65412:60
```

Supported Release 10.3.0E or later

ip prefix-list description

Configures a description of an IP prefix list.

Syntax `ip prefix-list name description`

Parameters

- name — Enter the name of the prefix list.
- *description* — Enter the description for the named prefix list.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information	The `no` version of this command removes the specified prefix list.
Example	`OS10(config)# ip prefix-list TEST description TEST_LIST`
Supported Release	10.3.0E or later

ip prefix-list deny

Creates a prefix list to deny route filtering from a specified network address.

Syntax	`ip prefix-list` *`name`* `deny [A.B.C.D/x [ge | le]]` *`prefix-len`*
Parameters	• *`name`* — Enter the name of the prefix list. • `A.B.C.D/x` — (Optional) Enter the source network address and mask in /prefix format (/x). • `ge` — Enter to indicate the network address is greater than or equal to the range specified. • `le` — Enter to indicate the network address is less than or equal to the range specified. • *`prefix-len`* — Enter the prefix length.
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command removes the specified prefix-list.
Example	`OS10(config)# ip prefix-list denyprefix deny 10.10.10.2/16 le 30`
Supported Release	10.3.0E or later

ip prefix-list permit

Creates a prefix-list to permit route filtering from a specified network address.

Syntax	`ip prefix-list` *`name`* `permit [A.B.C.D/x [ge | le]]` *`prefix-len`*
Parameters	• *`name`* — Enter the name of the prefix list. • `A.B.C.D/x` — (Optional) Enter the source network address and mask in /prefix format (/x). • `ge` — Enter to indicate the network address is greater than or equal to the range specified. • `le` — Enter to indicate the network address is less than or equal to the range specified. • *`prefix-len`* — Enter the prefix length.
Defaults	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command removes the specified prefix-list.
Example	`OS10(config)# ip prefix-list allowprefix permit 10.10.10.1/16 ge 10`
Supported Release	10.3.0E or later

ip prefix-list seq deny

Configures a filter to deny route filtering from a specified prefix list.

Syntax `ip prefix-list name seq num deny {A.B.C.D/x [ge | le] prefix-len}`

Parameters

- *name* — Enter the name of the prefix list.
- *num* — Enter the sequence list number.
- A.B.C.D/x — Enter the source network address and mask in /prefix format (/x).
- ge — Enter to indicate the network address is greater than or equal to the range specified.
- le — Enter to indicate the network address is less than or equal to the range specified.
- *prefix-len* — Enter the prefix length.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The no version of this command removes the specified prefix list.

Example

```
OS10(config)# ip prefix-list seqprefix seq 65535 deny 10.10.10.1/16 ge 10
```

Supported Release 10.3.0E or later

ip prefix-list seq permit

Configures a filter to permit route filtering from a specified prefix list.

Syntax `ipv6 prefix-list [name] seq num permit A::B/x [ge | le} prefix-len`

Parameters

- *name* — Enter the name of the prefix list.
- *num* — Enter the sequence list number.
- A.B.C.D/x — Enter the source network address and mask in /prefix format (/x).
- ge — Enter to indicate the network address is greater than or equal to the range specified.
- le — Enter to indicate the network address is less than or equal to the range specified.
- *prefix-len* — Enter the prefix length.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The no version of this command removes the specified prefix list.

Example

```
OS10(config)# ip prefix-list seqprefix seq 65535 permit 10.10.10.1/16 le 30
```

Supported Release 10.3.0E or later

ipv6 access-group

Configures an IPv6 access group.

Syntax `ipv6 access-group access-list-name {in | out}`

Parameters

- *access-list-name* — Enter the name of an IPv6 ACL. A maximum of 140 characters.
- `in` — Apply the ACL to incoming traffic.
- `out` — Apply the ACL to outgoing traffic.

Default Not configured

Command Mode INTERFACE

CONTROL-PLANE

Usage Information Use this command in the CONTROL-PLANE mode to apply a control-plane ACL. Control-plane ACLs are only applied on the ingress traffic. By default, the control-plane ACL is applied to the front-panel ports as well as the management port. The `no` version of this command deletes an IPv6 ACL configuration.

Example

```
OS10(conf-if-eth1/1/8)# ipv6 access-group test6 in
```

Example (Control-plane ACL)

```
OS10# configure terminal
OS10(config)# control-plane
OS10(config-control-plane)# ipv6 access-group aaa-cp-acl in
```

Supported Releases 10.2.0E or later; 10.4.1 or later (control-plane ACL)

ipv6 access-list

Creates an IP access list to filter based on an IPv6 address.

Syntax `ipv6 access-list access-list-name`

Parameters *access-list-name* — Enter the name of an IPv6 access list. A maximum of 140 characters.

Default Not configured

Command Mode CONFIGURATION

Usage Information None

Example

```
OS10(config)# ipv6 access-list acl6
```

Supported Release 10.2.0E or later

ipv6 prefix-list deny

Creates a prefix list to deny route filtering from a specified IPv6 network address.

Syntax `ipv6 prefix-list prefix-list-name deny {A::B/x [ge | le] prefix-len}`

Parameters

- *prefix-list-name* — Enter the IPv6 prefix list name.
- A::B/x — Enter the IPv6 address to deny.
- ge — Enter to indicate the network address is greater than or equal to the range specified.
- le — Enter to indicate the network address is less than or equal to the range specified.
- *prefix-len* — Enter the prefix length.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The no version of this command removes the specified prefix list.

Example

```
OS10(config)# ipv6 prefix-list TEST deny AB10::1/128 ge 10 le 30
```

Supported Release 10.3.0E or later

ipv6 prefix-list description

Configures a description of an IPv6 prefix-list.

Syntax `ipv6 prefix-list name description`

Parameters

- name — Enter the name of the IPv6 prefix-list.
- *description* — Enter the description for the named prefix-list.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The no version of this command removes the specified prefix list.

Example

```
OS10(config)# ipv6 prefix-list TEST description TEST_LIST
```

Supported Release 10.3.0E or later

ipv6 prefix-list permit

Creates a prefix-list to permit route filtering from a specified IPv6 network address.

Syntax `ipv6 prefix-list prefix-list-name permit {A::B/x [ge | le] prefix-len}`

Parameters

- *prefix-list-name* — Enter the IPv6 prefix-list name.
- A::B/x — Enter the IPv6 address to permit.
- ge — Enter to indicate the network address is greater than or equal to the range specified.
- le — Enter to indicate the network address is less than or equal to the range specified.
- *prefix-len* — Enter the prefix length.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The no version of this command removes the specified prefix-list.

Example

```
OS10(config)# ipv6 prefix-list TEST permit AB20::1/128 ge 10 le 30
```

Supported Release 10.3.0E or later

ipv6 prefix-list seq deny

Configures a filter to deny route filtering from a specified prefix-list.

Syntax

```
ipv6 prefix-list [name] seq num deny {A::B/x [ge | le] prefix-len}
```

Parameters

- `name` — (Optional) Enter the name of the IPv6 prefix-list.
- *`num`* — Enter the sequence number of the specified IPv6 prefix-list.
- `A::B/x` — Enter the IPv6 address and mask in /prefix format (/x).
- `ge` — Enter to indicate the network address is greater than or equal to the range specified.
- `le` — Enter to indicate the network address is less than or equal to the range specified.
- *`prefix-len`* — Enter the prefix length.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the specified prefix-list.

Example

```
OS10(config)# ipv6 prefix-list TEST seq 65535 deny AB20::1/128 ge 10
```

Supported Release 10.3.0E or later

ipv6 prefix-list seq permit

Configures a filter to permit route filtering from a specified prefix-list.

Syntax

```
ipv6 prefix-list [name] seq num permit A::B/x [ge | le} prefix-len
```

Parameters

- `name` — (Optional) Enter the name of the IPv6 prefix-list.
- *`num`* — Enter the sequence number of the specified IPv6 prefix list.
- `A::B/x` — Enter the IPv6 address and mask in /prefix format (/x).
- `ge` — Enter to indicate the network address is greater than or equal to the range specified.
- `le` — Enter to indicate the network address is less than or equal to the range specified.
- *`prefix-len`* — Enter the prefix length.

Defaults Not configured

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the specified prefix-list.

Example

```
OS10(config)# ipv6 prefix-list TEST seq 65535 permit AB10::1/128 ge 30
```

Supported Release 10.3.0E or later

mac access-group

Configures a MAC access group.

Syntax `mac access-group access-list-name {in | out}`

Parameters

- *access-list-name* — Enter the name of a MAC access list. A maximum of 140 characters.
- `in` — Apply the ACL to incoming traffic.
- `out` — Apply the ACL to outgoing traffic.

Default Not configured

Command Mode CONFIGURATION

CONTROL-PLANE

Usage Information Use this command in the CONTROL-PLANE mode to apply a control-plane ACL. Control-plane ACLs are only applied on the ingress traffic. By default, the control-plane ACL is applied to the front-panel ports. The `no` version of this command resets the value to the default.

Example

```
OS10(config)# mac access-group maclist in
OS10(conf-mac-acl)#
```

Example (Control-plane ACL)

```
OS10# configure terminal
OS10(config)# control-plane
OS10(config-control-plane)# mac access-group maclist in
```

Supported Releases 10.2.0E or later; 10.4.1 or later (control-plane ACL)

mac access-list

Creates a MAC access list to filter based on a MAC address.

Syntax `mac access-list access-list-name`

Parameters *access-list-name* — Enter the name of a MAC access list. A maximum of 140 characters.

Default Not configured

Command Mode CONFIGURATION

Usage Information None

Example

```
OS10(config)# mac access-list maclist
```

Supported Releases 10.2.0E or later

permit

Configures a filter to allow packets with a specific IPv4 address.

Syntax

```
permit [protocol-number | icmp | ip | tcp | udp] [A.B.C.D | A.B.C.D/x | any | host ip-address] [A.B.C.D | A.B.C.D/x | any | host ip-address] [capture | dscp value| fragment]
```

Parameters

- *protocol-number* — (Optional) Enter the protocol number identified in the IP header, from 0 to 255.
- icmp — (Optional) Enter the ICMP address to permit.
- ip — (Optional) Enter the IPv4 address to permit.
- tcp — (Optional) Enter the TCP address to permit.
- udp — (Optional) Enter the UDP address to permit.
- A.B.C.D — Enter the IPv4 address in dotted decimal format.
- A.B.C.D/x — Enter the number of bits that must match the dotted decimal address.
- any — (Optional) Set all routes which are subject to the filter:
  - capture — (Optional) Capture packets the filter processes.
  - dscp *value* — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
  - fragment — (Optional) Use ACLs to control packet fragments.
- host *ip-address* — (Optional) Enter the IPv4 address to use a host address only.

Default Not configured

Command Mode IPV4-ACL

Usage Information The no version of this command removes the filter.

Example

```
OS10(config)# ip access-list testflow
OS10(conf-ipv4-acl)# permit udp any any capture session 1
```

Supported Releases 10.2.0E or later

permit (IPv6)

Configures a filter to allow packets with a specific IPv6 address.

Syntax

```
permit [protocol-number | icmp | ipv6 | tcp | udp] [A::B | A::B/x | any | host ipv6-address] [A::B | A:B/x | any | host ipv6-address] [capture | dscp value | fragment]
```

Parameters

- *protocol-number* — (Optional) Enter the protocol number identified in the IPv6 header, from 0 to 255.
- icmp — (Optional) Enter the ICMP address to permit.
- ipv6 — (Optional) Enter the IPv6 address to permit.
- tcp — (Optional) Enter the TCP address to permit.
- udp — (Optional) Enter the UDP address to permit.
- A::B — Enter the IPv6 address in hexadecimal format separated by colons.
- A::B/x — Enter the number of bits that must match the IPv6 address.
- any — (Optional) Set all routes which are subject to the filter:

- capture — (Optional) Capture packets the filter processes.
- dscp *value* — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
- fragment — (Optional) Use ACLs to control packet fragments.

- host *ip-address* — (Optional) Enter the IPv6 address to use a host address only.

Default Not configured

Command Mode IPV6-ACL

Usage Information The no version of this command removes the filter.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# permit udp any any capture session 1
```

Supported Releases 10.2.0E or later

permit (MAC)

Configures a filter to allow packets with a specific MAC address.

Syntax

```
permit {nn:nn:nn:nn:nn:nn [00:00:00:00:00:00] | any} {nn:nn:nn:nn:nn:nn
[00:00:00:00:00:00] | any} [protocol-number | capture | cos | vlan]
```

Parameters

- *nn:nn:nn:nn:nn:nn* — Enter the MAC address.
- *00:00:00:00:00:00* — (Optional) Enter which bits in the MAC address must match. If you do not enter a mask, a mask of 00:00:00:00:00:00 applies.
- any — (Optional) Set which routes are subject to the filter:
  - *protocol-number* — Enter the MAC protocol number identified in the MAC header, from 600 to ffff.
  - capture — (Optional) Enter the capture packets the filter processes.
  - cos — (Optional) Enter the CoS value, from 0 to 7.
  - vlan — (Optional) Enter the VLAN number, from 1 to 4093.

Default Not configured

Command Mode MAC-ACL

Usage Information The no version of this command removes the filter.

Example

```
OS10(config)# mac access-list macacl
OS10(conf-mac-acl)# permit 00:00:00:00:11:11 00:00:11:11:11:11 any cos 7
OS10(conf-mac-acl)# permit 00:00:00:00:11:11 00:00:11:11:11:11 any vlan 2
```

Supported Releases 10.2.0E or later

permit icmp

Configures a filter to permit all or specific ICMP messages.

Syntax

```
permit icmp [A.B.C.D | A.B.C.D/x | any | host ip-address] [[A.B.C.D | A.B.C.D/x
| any | host ip-address] [capture | dscp value | fragment]
```

Parameters

- A.B.C.D — Enter the IPv4 address in dotted decimal format.

- `A.B.C.D/x` — Enter the number of bits that must match the dotted decimal address.
- `any` — (Optional) Set all routes which are subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp value` — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
- `host ip-address` — (Optional) Enter the IPv4 address to use a host address only.

Default Not configured

Command Mode IPV4-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# ip access-list testflow
OS10(conf-ipv4-acl)# permit icmp any any capture session 1
```

Supported Releases 10.2.0E or later

permit icmp (IPv6)

Configures a filter to permit all or specific ICMP messages.

Syntax

```
permit icmp [A::B | A::B/x | any | host ipv6-address] [A::B | A:B/x | any |
host ipv6-address] [capture | dscp value | fragment]
```

Parameters

- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits that must match the IPv6 address.
- `any` — (Optional) Set all routes which are subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp value` — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
- `host ipv6-address` — (Optional) Enter the IPv6 address to use a host address only.

Default Not configured

Command Mode IPV6-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# permit icmp any any capture session 1
```

Supported Releases 10.2.0E or later

permit ip

Configures a filter to permit all or specific packets from an IPv4 address.

Syntax

```
permit ip [A.B.C.D | A.B.C.D/x | any | host ip-address] [[A.B.C.D | A.B.C.D/x |
any | host ip-address] [capture | dscp | fragments]
```

Parameters

- `A.B.C.D` — Enter the IPv4 address in dotted decimal format.

- `A.B.C.D/x` — Enter the number of bits to match to the dotted decimal address.
- `any` — (Optional) Set all routes which are subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp value` — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
  - `fragments` — (Optional) Use ACLs to control packet fragments.
- `host ip-address` — (Optional) Enter the IPv4 address to use a host address only.

Default Not configured

Command Mode IPV4-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(conf-ipv4-acl)# permit ip any any capture session 1
```

Supported Releases 10.2.0E or later

permit ipv6

Configures a filter to permit all or specific packets from an IPv6 address.

Syntax

```
permit ipv6 [A::B | A::B/x | any | host ipv6-address] [A::B | A:B/x | any |
host ipv6-address] [capture| dscp | fragment]
```

Parameters

- `A::B` — (Optional) Enter the source IPv6 address from which the packet was sent and the destination address.
- `A::B/x` — (Optional) Enter the source network mask in /prefix format (/x) and the destination mask.
- `any` — (Optional) Set all routes which are subject to the filter:
  - `capture` — (Optional) Enter to capture packets the filter processes.
  - `dscp value` — (Optional) Enter to deny a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Enter to use ACLs to control packet fragments.
- `host ipv6-address` — Enter the IPv6 address to use a host address only.

Default Not configured

Command Mode IPV6-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(conf-ipv6-acl)# permit ipv6 any any count capture session 1
```

Supported Releases 10.2.0E or later

permit tcp

Configures a filter to permit TCP packets meeting the filter criteria.

Syntax

```
permit tcp [A.B.C.D | A.B.C.D/x | any | host ip-address [operator]] [[A.B.C.D |
A.B.C.D/x | any | host ip-address [operator] ] [ack | fin | psh | rst | syn |
urg] [capture | dscp value | fragment]
```

Parameters

- `A.B.C.D` — Enter the IPv4 address in dotted decimal format.
- `A.B.C.D/x` — Enter the number of bits that must match the dotted decimal address.
- `any` — (Optional) Set all routes which are subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp` *`value`* — (Optional) Permit a packet based on the DSCP values, 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
  - `ack` — (Optional) Set the bit as acknowledgement.
  - `fin` — (Optional) Set the bit as finish—no more data from sender.
  - `psh` — (Optional) Set the bit as push.
  - `rst` — (Optional) Set the bit as reset.
  - `syn` — (Optional) Set the bit as synchronize.
  - `urg` — (Optional) Set the bit set as urgent.

  NOTE: The control-plane ACLs do not support the `any` parameter.
- *`operator`* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - `eq` — Equal to
  - `gt` — Greater than
  - `lt` — Lesser than
  - `neq` — Not equal to
  - `range` — Range of ports, including the specified port numbers.

  NOTE: The control-plane ACLs support only the `eq` operator.
- `host` *`ip-address`* — (Optional) Enter the IPv4 address to use a host address only.

Default Not configured

Command Mode IPV4–ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(conf-ipv4-acl)# permit tcp any any capture session 1
```

Supported Releases 10.2.0E or later

permit tcp (IPv6)

Configures a filter to permit TCP packets meeting the filter criteria.

Syntax

```
permit tcp [A::B | A::B/x | any | host ipv6-address [eq | lt | gt | neq |
range]] [A::B | A:B/x | any | host ipv6-address [eq | lt | gt | neq | range]]
[ack | fin | psh | rst | syn | urg] [capture | dscp value | fragment]
```

Parameters

- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits that must match the IPv6 address.
- `any` — (Optional) Set all routes which are subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp` *`value`* — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.

NOTE: The control-plane ACLs do not support the `any` parameter.

- `host ipv6-address` — (Optional) Enter the IPv6 address to use a host address only.

Default Not configured

Command Mode IPV6-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# permit tcp any any capture session 1
```

Supported Releases 10.2.0E or later

permit udp

Configures a filter that allows UDP packets meeting the filter criteria.

Syntax

```
permit udp [A.B.C.D | A.B.C.D/x | any | host ip-address [eq | lt | gt | neq |
range]] [[A.B.C.D | A.B.C.D/x | any | host ip-address [eq | lt | gt | neq |
range] ] [ack | fin | psh | rst | syn | urg] [capture | dscp value | fragment]
```

Parameters

- `A.B.C.D` — Enter the IPv4 address in dotted decimal format.
- `A.B.C.D/x` — Enter the number of bits that must match the dotted decimal address.
- `any` — (Optional) Set all routes which are subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp value` — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
  - `eq` — (Optional) Permit packets which are equal to.
  - `lt` — (Optional) Permit packets which are less than.
  - `gt` — (Optional) Permit packets which are greater than.
  - `neq` — (Optional) Permit packets which are not equal to.
  - `range` — (Optional) Permit packets with a specific source and destination address.
  - `ack` — (Optional) Set the bit as acknowledgement.
  - `fin` — (Optional) Set the bit as finish—no more data from sender.
  - `psh` — (Optional) Set the bit as push.
  - `rst` — (Optional) Set the bit as reset.
  - `syn` — (Optional) Set the bit as synchronize.
  - `urg` — (Optional) Set the bit set as urgent.

NOTE: The control-plane ACL supports only the `eq` operator.

- `host ip-address` — (Optional) Enter the IPv4 address to use a host address only.

Default Not configured

Command Mode IPV4-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(config)# ip access-list testflow
OS10(conf-ipv4-acl)# permit udp any any capture session 1
```

Supported Releases 10.2.0E or later

permit udp (IPv6)

Configures a filter to permit UDP packets meeting the filter criteria.

Syntax

```
permit udp [A::B | A::B/x | any | host ipv6-address [operator]] [A::B | A:B/x |
any | host ipv6-address [operator]] [ack | fin | psh | rst | syn | urg]
[capture | dscp value | fragment]
```

Parameters

- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits that must match the IPv6 address.
- `any` — (Optional) Enter for all routes to be subject to the filter:
  - `capture` — (Optional) Capture packets the filter processes.
  - `dscp value` — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
  - `fragment` — (Optional) Use ACLs to control packet fragments.
  - `ack` — (Optional) Set the bit as acknowledgement.
  - `fin` — (Optional) Set the bit as finish—no more data from sender.
  - `psh` — (Optional) Set the bit as push.
  - `rst` — (Optional) Set the bit as reset.
  - `syn` — (Optional) Set the bit as synchronize.
  - `urg` — (Optional) Set the bit set as urgent.

    NOTE: The control-plane ACL supports only the `eq` operator.
- *`operator`* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - `eq` — Equal to
  - `gt` — Greater than
  - `lt` — Lesser than
  - `neq` — Not equal to
  - `range` — Range of ports, including the specified port numbers.
- `host ipv6-address` — (Optional) Enter the keyword and the IPv6 address to use a host address only.

Default Not configured

Command Mode IPV6-ACL

Usage Information The `no` version of this command removes the filter.

Example

```
OS10(conf-ipv6-acl)# permit udp any any capture session 1 count
```

Supported Releases 10.2.0E or later

remark

Specifies an ACL entry description.

Syntax

```
remark description
```

Parameters *description* — Enter a description. A maximum of 80 charaters.

Default Not configured

Command Mode IPV4-ACL

Usage Information Configure up to 16777214 remarks for a given IPv4, IPv6, or MAC. The `no` version of the command removes the ACL entry description.

Supported Releases 10.2.0E or later

seq deny

Assigns a sequence number to deny IPv4 addresses while creating the filter.

Syntax

```
seq sequence-number deny [protocol-number | icmp | ip | tcp | udp] [A.B.C.D |
A.B.C.D/x | any | host ip-address] [A.B.C.D | A.B.C.D/x | any | host ip-
address] [capture | count | dscp value | fragment | log]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the ACL for editing and sequencing number, from 1 to 16777214.
- *protocol-number* — (Optional) Enter the protocol number, from 0 to 255.
- `icmp` — (Optional) Enter the ICMP address to deny.
- `ip` — (Optional) Enter the IPv4 address to deny.
- `tcp` — (Optional) Enter the TCP address to deny.
- `udp` — (Optional) Enter the UDP address to deny.
- `A.B.C.D` — (Optional) Enter the IPv4 address in dotted decimal format.
- `A.B.C.D/x` — (Optional) Enter the number of bits that must match the dotted decimal address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *ip-address* — (Optional) Enter the IPv4 address to use a host address only.
- `capture` — (Optional) Capture packets the filter processes.
- `count` — (Optional) Count packets the filter processes.
- `byte` — (Optional) Count bytes the filter processes.
- `dscp` *value* — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV4-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *sequence-number* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ip access-list testflow
OS10(conf-ipv4-acl)# seq 10 deny tcp any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq deny (IPv6)

Assigns a sequence number to deny IPv6 addresses while creating the filter.

Syntax

```
seq sequence-number deny [protocol-number icmp | ip | tcp | udp] [A::B | A::B/x
| any | host ipv6-address] [A::B | A::B/x | any | host ipv6-address] [capture |
count | dscp value | fragment | log]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- *protocol-number* — (Optional) Enter the protocol number, from 0 to 255.
- `icmp` — (Optional) Enter the ICMP address to deny.
- `ip` — (Optional) Enter the IPv6 address to deny.
- `tcp` — (Optional) Enter the TCP address to deny.
- `udp` — (Optional) Enter the UDP address to deny.
- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits that must match the IPv6 address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *ipv6-address* — (Optional) Enter to use an IPv6 host address only.
- `capture` — (Optional) Enter to capture packets the filter processes.
- `count` — (Optional) Enter to count packets the filter processes.
- `byte` — (Optional) Enter to count bytes the filter processes.
- `dscp` *value* — (Optional) Enter to deny a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Enter to use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV6-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *sequence-number* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# seq 5 deny ipv6 any any capture session 1 count log
```

Supported Releases 10.2.0E or later

seq deny (MAC)

Assigns a sequence number to a deny filter in a MAC access list while creating the filter.

Syntax

```
seq sequence-number deny {nn:nn:nn:nn:nn:nn [00:00:00:00:00:00] | any}
{nn:nn:nn:nn:nn:nn [00:00:00:00:00:00] | any} [protocol-number | capture | cos
| vlan]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- *nn:nn:nn:nn:nn:nn* — Enter the source MAC address.
- *00:00:00:00:00:00* — (Optional) Enter which bits in the MAC address must match. If you do not enter a mask, a mask of 00:00:00:00:00:00 applies.
- any — (Optional) Set all routes which are subject to the filter:
  - *protocol-number* — Protocol number identified in the MAC header, from 600 to ffff.
  - capture — (Optional) Capture packets the filter processes.
  - cos — (Optional) CoS value, from 0 to 7.
  - vlan — (Optional) VLAN number, from 1 to 4093.

Default Not configured

Command Mode CONFIG-MAC-ACL

Usage Information The no version of this command removes the filter, or use the no seq *sequence-number* command if you know the filter's sequence number.

Example

```
OS10(config)# mac access-list macacl
OS10(conf-mac-acl)# seq 10 deny 00:00:00:00:11:11 00:00:11:11:11:11 any cos 7
OS10(conf-mac-acl)# seq 20 deny 00:00:00:00:11:11 00:00:11:11:11:11 any vlan 2
```

Supported Releases 10.2.0E or later

seq deny icmp

Assigns a filter to deny ICMP messages while creating the filter.

Syntax

```
seq sequence-number deny icmp [A.B.C.D | A.B.C.D/x | any | host ip-address]
[A.B.C.D | A.B.C.D/x | any | host ip-address] [capture | count | dscp value |
fragment | log]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- A.B.C.D — Enter the IPv4 address in dotted decimal format.
- A.B.C.D/x — Enter the number of bits that must match the dotted decimal address.
- any — (Optional) Enter the keyword any to specify any source or destination IP address.
- host *ip-address* — (Optional) Enter the IPv4 address to use a host IP address only.
- capture — (Optional) Capture packets the filter processes.
- count — (Optional) Count packets the filter processes.
- byte — (Optional) Count bytes the filter processes.
- dscp *value* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
- fragment — (Optional) Use ACLs to control packet fragments.
- log — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV4-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *`sequence-number`* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ip access-list egress
OS10(conf-ipv4-acl)# seq 5 deny icmp any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq deny icmp (IPv6)

Assigns a sequence number to deny ICMP messages while creating the filter.

Syntax

```
seq sequence-number deny icmp [A::B | A::B/x | any | host ipv6-address] [A::B |
A::B/x | any | host ipv6-address] [capture | count | dscp value | fragment |
log]
```

Parameters

- *`sequence-number`* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits that must match the IPv6 address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *`ipv6-address`* — (Optional) Enter the IPv6 address to use a host address only.
- `capture` — (Optional) Capture packets the filter processes.
- `dscp` *`value`* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV6-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *`sequence-number`* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# seq 10 deny icmp any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq deny ip

Assigns a sequence number to deny IPv4 addresses while creating the filter.

Syntax

```
seq sequence-number deny ip [A.B.C.D | A.B.C.D/x | any | host ip-address]
[A.B.C.D | A.B.C.D/x | any | host ip-address] [capture | count | dscp value |
fragment | log]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- `A.B.C.D` — Enter the IPv4 address in dotted decimal format.
- `A.B.C.D/x` — Enter the number of bits that must match the dotted decimal address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *ip-address* — (Optional) Enter the IPv4 address to use a host address only.
- `capture` — (Optional) Capture packets the filter processes.
- `count` — (Optional) Count packets the filter processes.
- `byte` — (Optional) Count bytes the filter processes.
- `dscp` *value* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV4-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *sequence-number* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ip access-list egress
OS10(config-ipv4-acl)# seq 10 deny ip any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq deny ipv6

Assigns a filter to deny IPv6 addresses while creating the filter.

Syntax

```
seq sequence-number deny ip [A::B | A::B/x | any | host ipv6-address] [A::B |
A:B/x | any | host ipv6-address] [capture | count | dscp value | fragment |
log]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits that must match the IPv6 address.
- `any` —(Optional) Enter the keyword `any` to specify any source or destination address.

- `host ip-address` — (Optional) Enter the IPv6 address to use a host address only.
- `capture` — (Optional) Capture packets the filter processes.
- `count` — (Optional) Count packets the filter processes.
- `byte` — (Optional) Count bytes the filter processes.
- `dscp value` — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV6-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq sequence-number` command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# seq 10 deny ipv6 any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq deny tcp

Assigns a filter to deny TCP packets while creating the filter.

Syntax

```
seq sequence-number deny tcp [A.B.C.D | A.B.C.D/x | any | host ip-address
[operator]] [[A.B.C.D | A.B.C.D/x | any | host ip-address [operator] ] [ack |
fin | psh | rst | syn | urg] [capture | count | dscp value | fragment | log]
```

Parameters

- `sequence-number` — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- `A.B.C.D` — Enter the IPv4 address in dotted decimal format.
- `A.B.C.D/x` — Enter the number of bits that must match the dotted decimal address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host ip-address` — (Optional) Enter the IPv4 address to use a host address only.
- `ack` — (Optional) Set the bit as acknowledgement.
- `fin` — (Optional) Set the bit as finish—no more data from sender.
- `psh` — (Optional) Set the bit as push.
- `rst` — (Optional) Set the bit as reset.
- `syn` — (Optional) Set the bit as synchronize.
- `urg` — (Optional) Set the bit set as urgent.
- `capture` — (Optional) Capture packets the filter processes.
- `dscp value` — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.
- `operator` — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:

- eq — Equal to
- gt — Greater than
- lt — Lesser than
- neq — Not equal to
- range — Range of ports, including the specified port numbers.

Default Not configured

Command Mode IPV4-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *`sequence-number`* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ip access-list egress
OS10(conf-ipv4-acl)# seq 10 deny tcp any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq deny tcp (IPv6)

Assigns a filter to deny TCP packets while creating the filter.

Syntax

```
seq sequence-number deny tcp [A::B | A::B/x | any | host ipv6-address
[operator]] [A::B | A:B/x | any | host ipv6-address [operator]] [ack | fin |
psh | rst | syn | urg] [capture | count | dscp value | fragment | log]
```

Parameters

- *`sequence-number`* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits that must match the IPv6 address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *`ip-address`* — (Optional) Enter the IPv6 address to use a host address only.
- `ack` — (Optional) Set the bit as acknowledgement.
- `fin` — (Optional) Set the bit as finish—no more data from sender.
- `psh` — (Optional) Set the bit as push.
- `rst` — (Optional) Set the bit as reset.
- `syn` — (Optional) Set the bit as synchronize.
- `urg` — (Optional) Set the bit set as urgent.
- `capture` — (Optional) Capture packets the filter processes.
- `dscp` *`value`* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.
- *`operator`* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - eq — Equal to
  - gt — Greater than

- `lt` — Lesser than
- `neq` — Not equal to
- `range` — Range of ports, including the specified port numbers.

Default Not configured

Command Mode IPV6-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq sequence-number` command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# seq 10 deny tcp any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq deny udp

Assigns a filter to deny UDP packets while creating the filter.

Syntax

```
seq sequence-number deny udp [A.B.C.D | A.B.C.D/x | any | host ip-address
[operator]] [[A.B.C.D | A.B.C.D/x | any | host ip-address [operator] ] [ack |
fin | psh | rst | syn | urg] [capture | count | dscp value | fragment | log]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- `A.B.C.D` — Enter the IPv4 address in dotted decimal format.
- `A.B.C.D/x` — Enter the number of bits that must match the dotted decimal address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *ip-address* — (Optional) Enter the IPv4 address to use a host address only.
- `ack` — (Optional) Set the bit as acknowledgment.
- `fin` — (Optional) Set the bit as finish—no more data from sender.
- `psh` — (Optional) Set the bit as push.
- `rst` — (Optional) Set the bit as reset.
- `syn` — (Optional) Set the bit as synchronize.
- `urg` — (Optional) Set the bit set as urgent.
- `capture` — (Optional) Capture packets the filter processes.
- `dscp` *value* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.
- *operator* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - `eq` — Equal to
  - `gt` — Greater than
  - `lt` — Lesser than
  - `neq` — Not equal to

- range — Range of ports, including the specified port numbers.

Default Not configured

Command Mode IPV4-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq sequence-number` command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ip access-list egress
OS10(conf-ipv4-acl)# seq 10 deny udp any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq deny udp (IPv6)

Assigns a filter to deny UDP packets while creating the filter.

Syntax

```
seq sequence-number deny udp [A::B | A::B/x | any | host ipv6-address
[operator]] [A::B | A:B/x | any | host ipv6-address [operator]] [ack | fin |
psh | rst | syn | urg] [capture | count | dscp value | fragment | log]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- A::B — Enter the IPv6 address in hexadecimal format separated by colons.
- A::B/x — Enter the number of bits that must match the IPv6 address.
- any — (Optional) Enter the keyword any to specify any source or destination IP address.
- host *ipv6-address* — (Optional) Enter the IPv6 address to use a host address only.
- ack — (Optional) Set the bit as acknowledgment.
- fin — (Optional) Set the bit as finish—no more data from sender.
- psh — (Optional) Set the bit as push.
- rst — (Optional) Set the bit as reset.
- syn — (Optional) Set the bit as synchronize.
- urg — (Optional) Set the bit set as urgent.
- capture — (Optional) Capture packets the filter processes.
- dscp *value* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
- fragment — (Optional) Use ACLs to control packet fragments.
- log — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.
- *operator* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - eq — Equal to
  - gt — Greater than
  - lt — Lesser than
  - neq — Not equal to
  - range — Range of ports, including the specified port numbers.

Default Not configured

Command Mode IPV6-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *`sequence-number`* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# seq 10 deny udp any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq permit

Assigns a sequence number to permit packets while creating the filter.

Syntax

```
seq sequence-number permit [protocol-number A.B.C.D | A.B.C.D/x | any | host
ip-address] [A.B.C.D | A.B.C.D/x | any | host ip-address] [capture | count |
dscp value | fragment | log]
```

Parameters

- *`sequence-number`* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- *`protocol-number`* — (Optional) Enter the protocol number, from 0 to 255.
- `A.B.C.D` — Enter the IPv4 address in dotted decimal format.
- `A.B.C.D/x` — Enter the number of bits that must match the dotted decimal address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *`ip-address`* — (Optional) Enter the IPv4 address to use a host address only.
- `capture` — (Optional) Capture packets the filter processes.
- `count` — (Optional) Count packets the filter processes.
- `byte` — (Optional) Count bytes the filter processes.
- `dscp` *`value`* — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV4-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *`sequence-number`* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ip access-list testflow
OS10(conf-ipv4-acl)# seq 10 permit ip any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq permit (IPv6)

Assigns a sequence number to permit IPv6 packets, while creating a filter.

Syntax

```
seq sequence-number permit protocol-number [A::B | A::B/x | any | host ipv6-address] [A::B | A:B/x | any | host ipv6-address] [capture | count | dscp value | fragment | log]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- *protocol-number* — (Optional) Enter the protocol number, from 0 to 255.
- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits that must match the IPv6 address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *ipv6-address* — (Optional) Enter the IPv6 address to be used as the host address.
- `capture` — (Optional) Enter to capture packets the filter processes.
- `count` — (Optional) Enter to count packets the filter processes.
- `byte` — (Optional) Enter to count bytes the filter processes.
- `dscp` *value* — (Optional) Enter the DSCP value to permit a packet, from 0 to 63.
- `fragment` — (Optional) Enter to use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV6-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *sequence-number* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# seq 10 permit ipv6 any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq permit (MAC)

Assigns a sequence number to permit MAC addresses while creating a filter.

Syntax

```
seq sequence-number permit {nn:nn:nn:nn:nn:nn [00:00:00:00:00:00] | any} {nn:nn:nn:nn:nn:nn [00:00:00:00:00:00] | any} [protocol-number | capture | cos | vlan]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing, from 1 to 16777214.
- *nn:nn:nn:nn:nn:nn* — Enter the MAC address of the network from or to which the packets were sent.

- *00:00:00:00:00:00* — (Optional) Enter which bits in the MAC address must match. If you do not enter a mask, a mask of 00:00:00:00:00:00 applies.
- any — (Optional) Set all routes to be subject to the filter:
  - *protocol-number* — (Optional) Enter the protocol number identified in the MAC header, from 600 to ffff.
  - capture — (Optional) Enter the capture packets the filter processes.
  - cos — (Optional) Enter the CoS value, from 0 to 7.
  - vlan — (Optional) Enter the VLAN number, from 1 to 4093.

Default Not configured

Command Mode MAC-ACL

Usage Information The no version of this command removes the filter, or use the no seq *sequence-number* command if you know the filter's sequence number.

Example

```
OS10(config)# mac access-list macacl
OS10(conf-mac-acl)# seq 10 permit 00:00:00:00:11:11 00:00:11:11:11:11 any cos 7
OS10(conf-mac-acl)# seq 20 permit 00:00:00:00:11:11 00:00:11:11:11:11 any vlan
2
```

Supported Releases 10.2.0E or later

seq permit icmp

Assigns a sequence number to allow ICMP messages while creating the filter

Syntax

```
seq sequence-number permit icmp [A.B.C.D | A.B.C.D/x | any | host ip-address]
[A.B.C.D | A.B.C.D/x | any | host ip-address] [capture | count | dscp value |
fragment | log]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- A.B.C.D — Enter the IPv4 address in dotted decimal format.
- A.B.C.D/x — Enter the number of bits that must match the dotted decimal address.
- any — (Optional) Enter the keyword any to specify any source or destination IP address.
- host *ip-address* — (Optional) Enter the IPv4 address to use a host address only.
- capture — (Optional) Capture packets the filter processes.
- count — (Optional) Count packets the filter processes.
- byte — (Optional) Count bytes the filter processes.
- dscp *value* — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
- fragment — (Optional) Use ACLs to control packet fragments.
- log — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV4-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the count byte options, only bytes increment. The no version of this command removes the filter, or use the no seq *sequence-number* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ip access-list egress
OS10(conf-ipv4-acl)# seq 5 permit icmp any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq permit icmp (IPv6)

Assigns a sequence number to allow ICMP messages while creating the filter.

Syntax

```
seq sequence-number permit icmp [A::B | A::B/x | any | host ipv6-address] [A::B
| A:B/x | any | host ipv6-address] [capture | count | dscp value | fragment |
log]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits that must match the IPv6 address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *ipv6-address* — (Optional) Enter the IPv6 address to use a host address only.
- `capture` — (Optional) Capture packets the filter processes.
- `count` — (Optional) Count packets the filter processes.
- `byte` — (Optional) Count bytes the filter processes.
- `dscp` *value* — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV6-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *sequence-number* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ipv6 access-list ipv6test
OS10(conf-ipv6-acl)# seq 5 permit icmp any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq permit ip

Assigns a sequence number to allow packets while creating the filter.

Syntax

```
seq sequence-number permit ip [A.B.C.D | A.B.C.D/x | any | host ip-address]
[A.B.C.D | A.B.C.D/x | any | host ip-address] [capture | count | dscp value |
fragment | log]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- A.B.C.D — Enter the IPv4 address in dotted decimal format.
- A.B.C.D/x — Enter the number of bits that must match the dotted decimal address.
- any — (Optional) Enter the keyword any to specify any source or destination IP address.
- host *ip-address* — (Optional) Enter the IPv4 address to use a host address only.
- capture — (Optional) Capture packets the filter processes.
- count — (Optional) Count packets the filter processes.
- byte — (Optional) Count bytes the filter processes.
- dscp *value* — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
- fragment — (Optional) Use ACLs to control packet fragments.
- log — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV4-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the count byte options, only bytes increment. The no version of this command removes the filter, or use the no seq *sequence-number* command if you know the filter's sequence number.

The count, byte, and log options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ip access-list egress
OS10(conf-ipv4-acl)# seq 5 permit ip any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq permit ipv6

Assigns a sequence number to allow packets while creating the filter.

Syntax

```
seq sequence-number permit ipv6 [A::B | A::B/x | any | host ipv6-address] [A::B
| A:B/x | any | host ipv6-address] [capture | count | dscp value | fragment |
log]
```

Parameters

- *sequence-number* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- A::B — Enter the IPv6 address in hexadecimal format separated by colons.
- A::B/x — Enter the number of bits that must match the IPv6 address.
- any — (Optional) Enter the keyword any to specify any source or destination IP address.
- host *ipv6-address* — (Optional) Enter the IPv6 address to use a host address only.
- capture — (Optional) Capture packets the filter processes.
- count — (Optional) Count packets the filter processes.
- byte — (Optional) Count bytes the filter processes.
- dscp *value* — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
- fragment — (Optional) Use ACLs to control packet fragments.
- log — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV6-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *`sequence-number`* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ipv6 access-list egress
OS10(conf-ipv6-acl)# seq 5 permit ipv6 any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq permit tcp

Assigns a sequence number to allow TCP packets while creating the filter.

Syntax

```
seq sequence-number permit tcp [A.B.C.D | A.B.C.D/x | any | host ip-address
[operator]] [[A.B.C.D | A.B.C.D/x | any | host ip-address [operator] ] [ack |
fin | psh | rst | syn | urg] [capture | count | dscp value | fragment | log]
```

Parameters

- *`sequence-number`* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- `A.B.C.D` — Enter the IPv4 address in dotted decimal format.
- `A.B.C.D/x` — Enter the number of bits that must match the dotted decimal address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *`ip-address`* — (Optional) Enter the IPv4 address to use a host address only.
- *`operator`* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - `eq` — Equal to
  - `gt` — Greater than
  - `lt` — Lesser than
  - `neq` — Not equal to
  - `range` — Range of ports, including the specified port numbers.
- `ack` — (Optional) Set the bit as acknowledgment.
- `fin` — (Optional) Set the bit as finish—no more data from sender.
- `psh` — (Optional) Set the bit as push.
- `rst` — (Optional) Set the bit as reset.
- `syn` — (Optional) Set the bit as synchronize.
- `urg` — (Optional) Set the bit set as urgent.
- `capture` — (Optional) Capture packets the filter processes.
- `count` — (Optional) Count packets the filter processes.
- `byte` — (Optional) Count bytes the filter processes.
- `dscp` *`value`* — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV4-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *`sequence-number`* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ip access-list egress
OS10(conf-ipv4-acl)# seq 5 permit tcp any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq permit tcp (IPv6)

Assigns a sequence number to allow TCP IPv6 packets while creating the filter.

Syntax

```
seq sequence-number permit tcp [A::B | A::B/x | any | host ipv6-address
[operator]] [A::B | A:B/x | any | host ipv6-address [operator]] [ack | fin |
psh | rst | syn | urg] [capture | count | dscp value | fragment | log]
```

Parameters

- *`sequence-number`* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits that must match the IPv6 address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *`ipv6-address`* — (Optional) Enter the IPv6 address to use a host address only.
- *`operator`* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - `eq` — Equal to
  - `gt` — Greater than
  - `lt` — Lesser than
  - `neq` — Not equal to
  - `range` — Range of ports, including the specified port numbers.
- `ack` — (Optional) Set the bit as acknowledgment.
- `fin` — (Optional) Set the bit as finish—no more data from sender.
- `psh` — (Optional) Set the bit as push.
- `rst` — (Optional) Set the bit as reset.
- `syn` — (Optional) Set the bit as synchronize.
- `urg` — (Optional) Set the bit set as urgent.
- `capture` — (Optional) Capture packets the filter processes.
- `count` — (Optional) Count packets the filter processes.
- `byte` — (Optional) Count bytes the filter processes.
- `dscp` *`value`* — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV6-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *`sequence-number`* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ipv6 access-list egress
OS10(conf-ipv6-acl)# seq 5 permit tcp any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq permit udp

Assigns a sequence number to allow UDP packets while creating the filter.

Syntax

```
seq sequence-number permit udp [A.B.C.D | A.B.C.D/x | any | host ip-address
[operator]] [[A.B.C.D | A.B.C.D/x | any | host ip-address [operator] ] [ack |
fin | psh | rst | syn | urg] [capture | count | dscp value | fragment | log]
```

Parameters

- *`sequence-number`* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- `A.B.C.D` — Enter the IPv4 address in dotted decimal format.
- `A.B.C.D/x` — Enter the number of bits that must match the dotted decimal address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *`ip-address`* — (Optional) Enter the IPv4 address to use a host address only.
- *`operator`* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - `eq` — Equal to
  - `gt` — Greater than
  - `lt` — Lesser than
  - `neq` — Not equal to
  - `range` — Range of ports, including the specified port numbers.
- `ack` — (Optional) Set the bit as acknowledgment.
- `fin` — (Optional) Set the bit as finish—no more data from sender.
- `psh` — (Optional) Set the bit as push.
- `rst` — (Optional) Set the bit as reset.
- `syn` — (Optional) Set the bit as synchronize.
- `urg` — (Optional) Set the bit set as urgent.
- `capture` — (Optional) Capture packets the filter processes.
- `count` — (Optional) Count packets the filter processes.
- `byte` — (Optional) Count bytes the filter processes.
- `dscp` *`value`* — (Optional) Deny a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV4-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *`sequence-number`* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ip access-list egress
OS10(conf-ipv4-acl)# seq 5 permit udp any any capture session 1 log
```

Supported Releases 10.2.0E or later

seq permit udp (IPv6)

Assigns a sequence number to allow UDP IPv6 packets while creating a filter.

Syntax

```
seq sequence-number permit udp [A::B | A::B/x | any | host ipv6-address
[operator]] [A::B | A:B/x | any | host ipv6-address [operator]] [ack | fin |
psh | rst | syn | urg] [capture | count | dscp value | fragment | log]
```

Parameters

- *`sequence-number`* — Enter the sequence number to identify the route-map for editing and sequencing number, from 1 to 16777214.
- `A::B` — Enter the IPv6 address in hexadecimal format separated by colons.
- `A::B/x` — Enter the number of bits that must match the IPv6 address.
- `any` — (Optional) Enter the keyword `any` to specify any source or destination IP address.
- `host` *`ipv6-address`* — (Optional) Enter the IPv6 address to use a host address only.
- *`operator`* — (Optional) Enter a logical operator to match the packets on the specified port number. The following options are available:
  - `eq` — Equal to
  - `gt` — Greater than
  - `lt` — Lesser than
  - `neq` — Not equal to
  - `range` — Range of ports, including the specified port numbers.
- `ack` — (Optional) Set the bit as acknowledgment.
- `fin` — (Optional) Set the bit as finish—no more data from sender.
- `psh` — (Optional) Set the bit as push.
- `rst` — (Optional) Set the bit as reset.
- `syn` — (Optional) Set the bit as synchronize.
- `urg` — (Optional) Set the bit set as urgent.
- `capture` — (Optional) Capture packets the filter processes.
- `count` — (Optional) Count packets the filter processes.
- `byte` — (Optional) Count bytes the filter processes.
- `dscp` *`value`* — (Optional) Permit a packet based on the DSCP values, from 0 to 63.
- `fragment` — (Optional) Use ACLs to control packet fragments.
- `log` — (Optional) Enables ACL logging. Information about packets that match an ACL rule are logged.

Default Not configured

Command Mode IPV6-ACL

Usage Information OS10 cannot count both packets and bytes; when you enter the `count byte` options, only bytes increment. The `no` version of this command removes the filter, or use the `no seq` *`sequence-number`* command if you know the filter's sequence number.

The `count`, `byte`, and `log` options are not supported on the S5148F-ON platform.

Example

```
OS10(config)# ipv6 access-list egress
OS10(conf-ipv6-acl)# seq 5 permit udp any any capture session 1 log
```

Supported Releases 10.2.0E or later

show access-group

Displays IP, MAC, or IPv6 access-group information.

Syntax

```
show {ip | mac | ipv6} access-group name
```

Parameters

- `ip` — View IP access group information.
- `mac` — View MAC access group information.
- `ipv6` — View IPv6 access group information.
- `access-group` *`name`* — Enter the name of the access group.

Default Not configured

Command Mode EXEC

Usage Information None

Example (IP)

```
OS10# show ip access-group aaa
Ingress IP access list aaa on ethernet1/1/1
Ingress IP access list aaa on ethernet1/1/2
Egress IP access list aaa on ethernet1/1/2
```

Example (MAC)

```
OS10# show mac access-group bbb
Ingress MAC access list aaa on ethernet1/1/1
Ingress MAC access list aaa on ethernet1/1/2
Egress MAC access list aaa on ethernet1/1/2
```

Example (IPv6)

```
OS10# show ipv6 access-group ccc
Ingress IPV6 access list aaa on ethernet1/1/1
Ingress IPV6 access list aaa on ethernet1/1/2
Egress IPV6 access list aaa on ethernet1/1/2
```

Example (Control-plane ACL - IP)

```
OS10# show ip access-group aaa-cp-acl
Ingress IP access-list aaa-cp-acl on control-plane data mgmt
```

Example (Control-plane ACL - MAC)

```
OS10# show mac access-group aaa-cp-acl
Ingress MAC access-list aaa-cp-acl on control-plane data
```

Example (Control-plane ACL - IPv6)

```
OS10# show ipv6 access-group aaa-cp-acl
Ingress IPV6 access-list aaa-cp-acl on control-plane data mgmt
```

Supported Releases 10.2.0E or later; 10.4.1 or later (control-plane ACL)

show access-lists

Displays IP, MAC, or IPv6 access-list information.

Syntax `show {ip | mac | ipv6} access-lists {in | out} access-list-name`

Parameters

- `ip` — View IP access list information.
- `mac` — View MAC access list information.
- `ipv6` — View IPv6 access list information.
- `access-lists in | out` — Enter either `access lists in` or `access lists out`.
- *`access-list–name`* — Enter the name of the access-list.

Default Not configured

Command Mode EXEC

Usage Information None

Example (MAC In)

```
OS10# show mac access-lists in
Ingress MAC access list aaa
Active on interfaces :
  ethernet1/1/1
  ethernet1/1/2
 seq 10 permit any any
 seq 20 permit 11:11:11:11:11:11 22:22:22:22:22:22 any monitor
```

Example (MAC Out)

```
OS10# show mac access-lists out
Egress MAC access list aaa
 Active on interfaces :
  ethernet1/1/1
  ethernet1/1/2
 seq 10 permit any any
 seq 20 permit 11:11:11:11:11:11 22:22:22:22:22:22 any monitor
```

Example (IP In)

```
OS10# show ip access-lists in
Ingress IP access list aaaa
 Active on interfaces :
  ethernet1/1/1
  ethernet1/1/2
 seq 10 permit ip any any log
 seq 20 permit tcp any any
 seq 30 permit udp any any
```

Example (IP Out)

```
OS10# show ip access-lists out
Egress IP access list aaaa
 Active on interfaces :
  ethernet1/1/1
  ethernet1/1/2
 seq 10 permit ip any any
 seq 20 permit tcp any any
 seq 30 permit udp any any
```

Example (IPv6 In)

```
OS10# show ipv6 access-lists in
Ingress IPV6 access list bbb
 Active on interfaces :
  ethernet1/1/1
  ethernet1/1/2
```

```
 seq 10 permit any any
Ingress IPV6 access list ggg
 Active on interfaces :
  ethernet 1/1/3
 seq 5 permit ipv6 11::/32 any log
```

Example (IPv6 Out)

```
OS10# show ipv6 access-lists out
Egress IPV6 access list bbb
 Active on interfaces :
  ethernet1/1/1
  ethernet1/1/2
 seq 10 permit any any
Egress IPV6 access list ggg
 Active on interfaces :
  ethernet 1/1/1
 seq 5 permit ipv6 11::/32 any
```

Example (IP In - Control-plane ACL)

```
OS10# show ip access-lists in
Ingress IP access-list aaa-cp-acl
 Active on interfaces :
  control-plane data
 seq 10 permit ip any any
  control-plane mgmt
 seq 10 permit ip any any
```

Example (IPv6 In - Control-plane ACL)

```
OS10# show ipv6 access-lists in
Ingress IPV6 access-list aaa-cp-acl
 Active on interfaces :
  control-plane data
 seq 10 permit ipv6 any any
  control-plane mgmt
 seq 10 permit ipv6 any any
```

Example (MAC In - Control-plane ACL)

```
OS10# show mac access-lists in
Ingress MAC access-list mac-cp1
 Active on interfaces :
control-plane data
 seq 10 deny any any count (159 packets)
```

Supported Releases 10.2.0E or later; 10.4.1 or later (control-plane ACL)

show acl-table-usage detail

Displays the ingress and egress ACL tables, the features that are used, and their space utilizations.

Syntax `show acl-table-usage detail`

Parameters None

Default None

Command Mode EXEC

Usage Information The hardware pool displays the ingress application groups (pools), the features mapped to each of these groups, and the amou space available in each of the pools. The amount of space required to store a single ACL rule in a pool depends on the keywidt

The service pool displays the amount of used and free space for each of the features. The number of ACL rules configured for in the configured rules column. The number of used rows depends on the number of ports the configured rules are applied on.

Examples

Z9100-ON platform

```
OS10# show acl-table-usage detail

Ingress ACL utilization - Pipe 0
Hardware Pools
-------------------------------------------------------------------
Pool ID     App(s)            Used rows    Free rows    Max rows
-------------------------------------------------------------------
0         SYSTEM_FLOW          98           414          512
1         SYSTEM_FLOW          98           414          512
2         SYSTEM_FLOW          98           414          512
3         USER_IPV4_ACL        4            508          512
4         USER_IPV4_ACL        4            508          512
5         FREE                 0            512          512
6         USER_IPV6_ACL        4            508          512
7         USER_IPV6_ACL        4            508          512
8         USER_IPV6_ACL        4            508          512
9         USER_L2_ACL          4            508          512
10        USER_L2_ACL          4            508          512
11        FREE                 0            512          512
-------------------------------------------------------------------------------------------
Service Pools
-------------------------------------------------------------------------------------------
App                 Allocated pools  App group   Configured rules    Used rows    Free
-------------------------------------------------------------------------------------------
USER_L2_ACL         Shared:2         G9          1                   2            254
256
USER_IPV4_ACL       Shared:2         G3          1                   2            254
256
USER_IPV6_ACL       Shared:3         G6          1                   2            254
256
SYSTEM_FLOW         Shared:3         G0          49                  49           207
256
-------------------------------------------------------------------------------------------
Ingress ACL utilization - Pipe 1
Hardware Pools
-------------------------------------------------------------------
Pool ID     App(s)            Used rows    Free rows    Max rows
-------------------------------------------------------------------
0         SYSTEM_FLOW          98           414          512
1         SYSTEM_FLOW          98           414          512
2         SYSTEM_FLOW          98           414          512
3         USER_IPV4_ACL        0            512          512
4         USER_IPV4_ACL        0            512          512
5         FREE                 0            512          512
6         USER_IPV6_ACL        0            512          512
7         USER_IPV6_ACL        0            512          512
8         USER_IPV6_ACL        0            512          512
9         USER_L2_ACL          0            512          512
10        USER_L2_ACL          0            512          512
11        FREE                 0            512          512
-------------------------------------------------------------------------------------------
Service Pools
-------------------------------------------------------------------------------------------
App                 Allocated pools  App group   Configured rules    Used rows    Free
-------------------------------------------------------------------------------------------
SYSTEM_FLOW         Shared:3         G0          49                  49           207
256
-------------------------------------------------------------------------------------------
Ingress ACL utilization - Pipe 2
Hardware Pools
-------------------------------------------------------------------
Pool ID     App(s)            Used rows    Free rows    Max rows
-------------------------------------------------------------------
0         SYSTEM_FLOW          98           414          512
1         SYSTEM_FLOW          98           414          512
2         SYSTEM_FLOW          98           414          512
```

```
3         USER_IPV4_ACL       0            512          512
4         USER_IPV4_ACL       0            512          512
5         FREE                0            512          512
6         USER_IPV6_ACL       0            512          512
7         USER_IPV6_ACL       0            512          512
8         USER_IPV6_ACL       0            512          512
9         USER_L2_ACL         0            512          512
10        USER_L2_ACL         0            512          512
11        FREE                0            512          512
-------------------------------------------------------------------------------------------
Service Pools
-------------------------------------------------------------------------------------------
App                 Allocated pools  App group   Configured rules    Used rows    Free
-------------------------------------------------------------------------------------------
SYSTEM_FLOW         Shared:3         G0          49                  49           207
256
-------------------------------------------------------------------------------------------
Ingress ACL utilization - Pipe 3
Hardware Pools
-----------------------------------------------------------------------
Pool ID     App(s)               Used rows    Free rows    Max rows
-----------------------------------------------------------------------
0         SYSTEM_FLOW             98           414          512
1         SYSTEM_FLOW             98           414          512
2         SYSTEM_FLOW             98           414          512
3         USER_IPV4_ACL           0            512          512
4         USER_IPV4_ACL           0            512          512
5         FREE                    0            512          512
6         USER_IPV6_ACL           0            512          512
7         USER_IPV6_ACL           0            512          512
8         USER_IPV6_ACL           0            512          512
9         USER_L2_ACL             0            512          512
10        USER_L2_ACL             0            512          512
11        FREE                    0            512          512
-------------------------------------------------------------------------------------------
Service Pools
-------------------------------------------------------------------------------------------
App                 Allocated pools  App group   Configured rules    Used rows    Free
-------------------------------------------------------------------------------------------
SYSTEM_FLOW         Shared:3         G0          49                  49           207
256
-------------------------------------------------------------------------------------------
Egress ACL utilization
Hardware Pools
-------------------------------------------------------------------------------------------
Pool ID     App(s)                                                    Used rows    Free
-------------------------------------------------------------------------------------------
0         FREE                                                        0            256
256
1         FREE                                                        0            256
256
2         FREE                                                        0            256
256
3         FREE                                                        0            256
256
-------------------------------------------------------------------------------------------
Service Pools
-------------------------------------------------------------------------------------------
App                 Allocated pools  App group   Configured rules    Used rows    Free
-------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------
```

S6010-ON platform

```
OS10# show acl-table-usage detail
Ingress ACL utilization
Hardware Pools
-----------------------------------------------------------------------
```

```
Pool ID     App(s)             Used rows    Free rows    Max rows
-------------------------------------------------------------------
0          SYSTEM_FLOW          49           975          1024
1          SYSTEM_FLOW          49           975          1024
2          USER_IPV4_ACL        3            1021         1024
3          USER_L2_ACL          2            1022         1024
4          USER_IPV6_ACL        2            510          512
5          USER_IPV6_ACL        2            510          512
6          FCOE                 55           457          512
7          FCOE                 55           457          512
8          ISCSI_SNOOPING       12           500          512
9          FREE                 0            512          512
10         PBR_V6               1            511          512
11         PBR_V6               1            511          512
-------------------------------------------------------------------------------------------
Service Pools
-------------------------------------------------------------------------------------------
App                 Allocated pools  App group   Configured rules    Used rows    Free
-------------------------------------------------------------------------------------------
USER_L2_ACL         Shared:1         G3          1                   2            1022
1024
USER_IPV4_ACL       Shared:1         G2          2                   3            1021
1024
USER_IPV6_ACL       Shared:2         G4          1                   2            510
512
PBR_V6              Shared:2         G10         1                   1            511
512
SYSTEM_FLOW         Shared:2         G0          49                  49           975
1024
ISCSI_SNOOPING      Shared:1         G8          12                  12           500
512
FCOE                Shared:2         G6          55                  55           457
512
-------------------------------------------------------------------------------------------
Egress ACL utilization
Hardware Pools
-------------------------------------------------------------------
Pool ID     App(s)             Used rows    Free rows    Max rows
-------------------------------------------------------------------
0          USER_IPV4_EGRESS      2            254          256
1          USER_L2_ACL_EGRESS    2            254          256
2          USER_IPV6_EGRESS      2            254          256
3          USER_IPV6_EGRESS      2            254          256
-------------------------------------------------------------------------------------------
Service Pools
-------------------------------------------------------------------------------------------
App                 Allocated pools  App group   Configured rules    Used rows    Free
-------------------------------------------------------------------------------------------
USER_L2_ACL_EGRESS  Shared:1         G1          1                   2            254
256
USER_IPV4_EGRESS    Shared:1         G0          1                   2            254
256
USER_IPV6_EGRESS    Shared:2         G2          1                   2            254
```

Supported Releases 10.4.2 and later

show ip as-path-access-list

Displays the configured AS path access lists.

Syntax `show ip as-path-access-list [name]`

Parameters *name* — (Optional) Specify the name of the AS path access list.

Defaults None

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip as-path-access-list
ip as-path access-list hello
        permit 123
        deny 35
```

Supported Releases 10.3.0E or later

show ip community-list

Displays the configured IP community lists in alphabetic order.

Syntax `show ip community-list [name]`

Parameters *name* — (Optional) Enter the name of the standard IP community list. A maximum of 140 characters.

Defaults None

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip community-list
Standard Community List hello
        deny local-AS
        permit no-export
        deny 1:1
```

Supported Releases 10.3.0E or later

show ip extcommunity-list

Displays the configured IP external community lists in alphabetic order.

Syntax `show ip extcommunity-list [name]`

Parameters *name* — (Optional) Enter the name of the extended IP external community list. A maximum of 140 characters.

Defaults None

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip extcommunity-list
Standard Extended Community List hello
        permit RT:1:1
        deny SOO:1:4
```

Supported Releases 10.3.0E or later

show ip prefix-list

Displays configured IPv4 or IPv6 prefix list information.

Syntax

```
show {ip | ipv6} prefix-list [prefix-name]
```

Parameters

- `ip | ipv6`—(Optional) Displays information related to IPv4 or IPv6.
- *`prefix-name`* — Enter a text string for the prefix list name. A maximum of 140 characters.

Defaults None

Command Mode EXEC

Usage Information None

Example

```
OS10# show ip prefix-list
ip prefix-list hello:
seq 10 deny 1.2.3.4/24
seq 20 permit 3.4.4.5/32
```

Example (IPv6)

```
OS10# show ipv6 prefix-list
ipv6 prefix-list hello:
seq 10 permit 1::1/64
seq 20 deny 2::2/64
```

Supported Releases 10.3.0E or later

show logging access-list

Displays the ACL logging threshold and interval configuration.

Syntax

```
show logging access-list
```

Parameters None

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show logging access-list
ACL Logging
Threshold      :   10
Interval       :   5
```

Supported Releases 10.4.3.0 or later

Route-map commands

continue

Configures the next sequence of the route map.

Syntax `continue seq-number`

Parameters *seq-number* — Enter the next sequence number, from 1 to 65535.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes a match.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# continue 65535
```

Supported Releases 10.3.0E or later

match as-path

Configures a filter to match routes that have a certain AS path in their BGP paths.

Syntax `match as-path as-path-name`

Parameters *as-path-name* — Enter the name of an established AS-PATH ACL. A maximum of 140 characters.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes a match AS path filter.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# match as-path pathtest1
```

Supported Releases 10.3.0E or later

match community

Configures a filter to match routes that have a certain COMMUNITY attribute in their BGP path.

Syntax `match community community-list-name [exact-match]`

Parameters

- *community-list-name* — Enter the name of a configured community list.
- `exact-match` — (Optional) Select only those routes with the specified community list name.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes the community match filter.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# match community commlist1 exact-match
```

Supported Releases 10.3.0E or later

match extcommunity

Configures a filter to match routes that have a certain EXTCOMMUNITY attribute in their BGP path.

Syntax

```
match extcommunity extcommunity-list-name [exact-match]
```

Parameters

- *extcommunity-list-name* — Enter the name of a configured extcommunity list.
- `exact-match` — (Optional) Select only those routes with the specified extcommunity list name.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes the extcommunity match filter.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# match extcommunity extcommlist1 exact-match
```

Supported Releases 10.3.0E or later

match interface

Configures a filter to match routes whose next-hop is the configured interface.

Syntax

```
match interface interface
```

Parameters *interface* — Interface type:

- `ethernet` *node/slot/port[:subport]* — Enter the Ethernet interface information as the next-hop interface.
- `port-channel` *id-number* — Enter the port-channel number as the next-hop interface, from 1 to 128.
- `vlan` *vlan-id* —Enter the VLAN number as the next-hop interface, from 1 to 4093.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes the match.

Example

```
OS10(conf-route-map)# match interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)#
```

Supported Releases 10.2.0E or later

match ip address

Configures a filter to match routes based on IP addresses specified in IP prefix lists.

Syntax

```
match ip address {prefix-list prefix-list-name | access-list-name}
```

Parameters

- *prefix-list-name* — Enter the name of the configured prefix list. A maximum of 140 characters.

- *access-list-name* — Enter the name of the configured access list.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes a match.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# match ip address  prefix-list test10
```

Supported Releases 10.3.0E or later

match ip next-hop

Configures a filter to match based on the next-hop IP addresses specified in IP prefix lists.

Syntax `match ip next-hop prefix-list prefix-list`

Parameters *prefix-list* — Enter the name of the configured prefix list. A maximum of 140 characters.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes the match.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# match ip next-hop  prefix-list test100
```

Supported Releases 10.3.0E or later

match ipv6 address

Configures a filter to match routes based on IPv6 addresses specified in IP prefix lists.

Syntax `match ipv6 address {prefix-list prefix-list | access-list}`

Parameters

- *prefix-list* — Enter the name of the configured prefix list. A maximum of 140 characters.
- *access-list* — Enter the name of the access group or list.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes the match.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# match ipv6 address test100
```

Supported Releases 10.3.0E or later

match ipv6 next-hop

Configures a filter to match based on the next-hop IPv6 addresses specified in IP prefix lists.

Syntax `match ipv6 next-hop prefix-list prefix-list`

Parameters *prefix-list* — Enter the name of the configured prefix list. A maximum of 140 characters.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes the match.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# match ipv6 next-hop  prefix-list test100
```

Supported Releases 10.3.0E or later

match metric

Configures a filter to match on a specific value.

Syntax `match metric metric-value`

Parameters *metric-value* — Enter a value to match the route metric against, from 0 to 4294967295.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes the match.

Example

```
OS10(conf-route-map)# match metric 429132
```

Supported Releases 10.2.0E or later

match origin

Configures a filter to match routes based on the origin attribute of BGP.

Syntax `match origin {egp | igp | incomplete}`

Parameters

- `egp` — Match only remote EGP routes.
- `igp` — Match only on local IGP routes.
- `incomplete` — Match on unknown routes that are learned through some other means.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes the match.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# match origin egp
```

Supported Releases 10.3.0E or later

match route-type

Configures a filter to match routes based on how the route is defined.

Syntax

```
match route-type {{external {type-1 | type-2} | internal | local }
```

Parameters

- `external` — Match only on external OSPF routes. Enter the keyword then one of the following:
  - `type-1` — Match only on OSPF Type 1 routes.
  - `type-2` — Match only on OSPF Type 2 routes.
-
- `internal` — Match only on routes generated within OSPF areas.
- `local` — Match only on routes generated locally.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes the match.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# match route-type external type-1
```

Supported Releases 10.3.0E or later

match tag

Configures a filter to redistribute only routes that match a specific tag value.

Syntax

```
match tag tag-value
```

Parameters *tag-value* — Enter the tag value to match with the tag number, from 0 to 4294967295.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes the match.

Example

```
OS10(conf-route-map)# match tag 656442
```

Supported Releases 10.2.0E or later

route-map

Enables a route-map statement and configures its action and sequence number.

Syntax

```
route-map map-name [permit | deny | sequence-number]
```

Parameters

- *map-name* — Enter the name of the route-map. A maximum of 140 characters.

- *sequence-number* — (Optional) Enter the number to identify the route-map for editing and sequencing number from 1 to 65535. The default is 10.
- `permit` — (Optional) Set the route-map default as permit.
- `deny` — (Optional) Set the route default as deny.

Default Not configured

Command Mode CONFIGURATION

Usage Information

NOTE: Exercise caution when you delete route-maps — if you do not enter a sequence number, all route-maps with the same map-name are deleted.

The `no` version of this command removes a route-map.

Example

```
OS10(config)# route-map route1 permit 100
OS10(config-route-map)#
```

Supported Releases 10.2.0E or later

set comm-list add

Add communities in the specified list to the COMMUNITY attribute in a matching inbound or outbound BGP route.

Syntax `set comm-list {community-list-name} add`

Parameters *community-list-name* — Enter the name of an established community list. A maximum of 140 characters.

Defaults None

Command Mode ROUTE-MAP

Usage Information In a route map, use this `set` command to add a list of communities that pass a permit statement to the COMMUNITY attribute of a BGP route sent or received from a BGP peer. Use the `set comm-list delete` command to delete a community list from a matching route.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# set comm-list comlist1 add
```

Supported Releases 10.4.0E(R1) or later

set comm-list delete

Remove communities in the specified list from the COMMUNITY attribute in a matching inbound or outbound BGP route.

Syntax `set comm-list {community-list-name} delete`

Parameters *community-list-name* — Enter the name of an established community list. A maximum of 140 characters.

Defaults None‘

Command Mode ROUTE-MAP

Usage Information Configure the community list you use in the `set comm-list delete` command so that each filter contains only one community. For example, the filter `deny 100:12` is acceptable, but the filter `deny 120:13 140:33` results in an error. If you configure the `set comm-list delete` command and the `set community` command in the same route map sequence, the deletion `set comm-list delete` command processes before

the insertion `set community` command . To add communities in a community list to the COMMUNITY attribute in a BGP route, use the `set comm-list add` command.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# set comm-list comlist1 delete
```

Supported Releases 10.3.0E or later

set community

Sets the community attribute in BGP updates.

Syntax `set community {none | community-number}`

Parameters

- `none` — Enter to remove the community attribute from routes meeting the route map criteria.
- *`community-number`* — Enter the community number in `aa:nn` format, where `aa` is the AS number, 2 bytes, and `nn` is a value specific to that AS.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes a BGP COMMUNITY attribute assignment.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# set community none
```

Supported Releases 10.3.0E or later

set extcomm-list add

Add communities in the specified list to the EXTCOMMUNITY attribute in a matching inbound or outbound BGP route.

Syntax `set extcomm-list` *`extcommunity-list-name`* `add`

Parameter *`extcommunity-list-name`* — Enter the name of an established extcommunity list. A maximum of 140 characters.

Defaults None

Command Mode ROUTE-MAP

Usage Information In a route map, use this `set` command to add an extended list of communities that pass a permit statement to the EXTCOMMUNITY attribute of a BGP route sent or received from a BGP peer. Use the `set extcomm-list delete` command to delete an extended community list from a matching route.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# set extcomm-list TestList add
```

Supported Releases 10.4.0E(R1) or later

set extcomm-list delete

Remove communities in the specified list from the EXTCOMMUNITY attribute in a matching inbound or outbound BGP route.

Syntax `set extcomm-list extcommunity-list-name delete`

Parameter *extcommunity-list-name* — Enter the name of an established extcommunity list. A maximum of 140 characters.

Defaults None

Command Mode ROUTE-MAP

Usage Information To add communities in an extcommunity list to the EXTCOMMUNITY attribute in a BGP route, use the `set extcomm-list add` command.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# set extcomm-list TestList delete
```

Supported Releases 10.3.0E or later

set extcommunity

Sets the extended community attributes in a route map for BGP updates.

Syntax `set extcommunity rt {asn2:nn | asn4:nnnn | ip-addr:nn}`

Parameters

- `asn2:nn` — Enter an AS number in 2-byte format; for example, 1–65535:1–4294967295.
- `asn4:nnnn` — Enter an AS number in 4-byte format; for example, 1–4294967295:1–65535 or 1–65535.1–65535:1–65535.
- `ip-addr:nn` — Enter an AS number in dotted format, from 1 to 65535.

Default Not configured

Command Mode ROUTE-MAP

Usage Information The `no` version of this command deletes the set clause from a route map.

Example

```
OS10(config)# route-map bgp
OS10(conf-route-map)# set extcommunity rt 10.10.10.2:325
```

Supported Releases 10.3.0E or later

set local-preference

Sets the preference value for the AS path.

Syntax `set local-preference value`

Parameters *value* — Enter a number as the LOCAL_PREF attribute value, from 0 to 4294967295.

Default Not configured

Command Mode ROUTE-MAP

Usage Information This command changes the LOCAL_PREF attribute for routes meeting the route map criteria. To change the LOCAL_PREF for all routes, use the `bgp default local-preference` command. The `no` version of this command removes the LOCAL_PREF attribute.

Example
```
OS10(conf-route-map)# set local-preference 200
```

Supported Releases 10.2.0E or later

set metric

Set a metric value for a routing protocol.

Syntax `set metric [+ | -] metric-value`

Parameters

- `+` — (Optional) Add a metric value to the redistributed routes.
- `-` — (Optional) Subtract a metric value from the redistributed routes.
- *`metric-value`* — Enter a new metric value, from 0 to 4294967295.

Default Not configured

Command Mode ROUTE-MAP

Usage Information To establish an absolute metric, do not enter a plus or minus sign before the metric value. To establish a relative metric, enter a plus or minus sign immediately preceding the metric value. The value is added to or subtracted from the metric of any routes matching the route map. You cannot use both an absolute metric and a relative metric within the same route map sequence. Setting either metric overrides any previously configured value. The `no` version of this command removes the filter.

Example (Absolute)
```
OS10(conf-route-map)# set metric 10
```

Example (Relative)
```
OS10(conf-route-map)# set metric -25
```

Supported Releases 10.2.0E or later

set metric-type

Set the metric type for the a redistributed routel.

Syntax `set metric-type {type-1 | type-2 | external}`

Parameters

- `type-1` — Adds a route to an existing community.
- `type-2` — Sends a route in the local AS.
- `external` — Disables advertisement to peers.

Default Not configured

Command Mode ROUTE-MAP

Usage Information

- **BGP**

Affects BGP behavior only in outbound route maps and has no effect on other types of route maps. If the route map contains both a `set metric-type` and a `set metric` clause, the `set metric` clause takes precedence. If you enter the `internal` metric type in a BGP outbound route map, BGP sets the MED of the advertised routes to the IGP cost of the next hop of the advertised route. If the cost of the next hop changes, BGP is not forced to readvertise the route.

  - `external` — Reverts to the normal BGP rules for propagating the MED, the default.
  - `internal` — Sets the MED of a received route that is being propagated to an external peer equal to the IGP costs of the indirect next hop.
- **OSPF**
  - `external` — Sets the cost of the external routes so that it is equal to the sum of all internal costs and the external cost.
  - `internal` — Sets the cost of the external routes so that it is equal to the external cost alone, the default.

The `no` version of this command removes the set clause from a route map.

Example

```
OS10(conf-route-map)# set metric-type internal
```

Supported Releases 10.2.0E or later

set next-hop

Sets an IPv4 or IPv6 address as the next-hop.

Syntax `set {ip | ipv6} next-hop ip-address`

Parameters *ip-address* — Enter the IPv4 or IPv6 address for the next-hop.

Default Not configured

Command Mode ROUTE-MAP

Usage Information If you apply a route-map with the `set next-hop` command in ROUTER-BGP mode, it takes precedence over the `next-hop-self` command used in ROUTER-NEIGHBOR mode. In a route-map configuration, to configure more than one next-hop entry, use multiple `set {ip | ipv6} next-hop` commands. When you apply a route-map for redistribution or route updates in ROUTER-BGP mode, configure only one next-hop. Configure multiple next-hop entries only in a route-map used for other features. The `no` version of this command deletes the setting.

Example

```
OS10(conf-route-map)# set ip next-hop 10.10.10.2
```

Example (IPv6)

```
OS10(conf-route-map)# set ipv6 next-hop 11AA:22CC::9
```

Supported Releases 10.2.0E or later

set origin

Set the origin of the advertised route.

Syntax `set origin {egp | igp | incomplete}`

Parameters

- `egp` — Enter to add to existing community.
- `igp` — Enter to send inside the local-AS.

- `incomplete` — Enter to not advertise to peers.

Default	Not configured
Command Mode	ROUTE-MAP
Usage Information	The `no` version of this command deletes the set clause from a route map.
Example	`OS10(conf-route-map)# set origin egp`
Supported Releases	10.2.0E or later

set tag

Sets a tag for redistributed routes.

Syntax	`set tag tag-value`
Parameters	*`tag-value`* — Enter a tag number for the route to redistribute, from 0 to 4294967295.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command deletes the set clause from a route map.
Example	`OS10(conf-route-map)# set tag 23`
Supported Releases	10.2.0E or later

set weight

Set the BGP weight for the routing table.

Syntax	`set weight weight`
Parameters	*`weight`* — Enter a number as the weight the route uses to meet the route map specification, from 0 to 65535.
Default	Default router-originated is 32768 — all other routes are 0.
Command Mode	ROUTE-MAP
Usage Information	The `no` version of the command deletes the set clause from the route map.
Example	`OS10(conf-route-map)# set weight 200`
Supported Releases	10.2.0E or later

show route-map

Displays the current route map configurations.

Syntax	`show route-map [map-name]`
Parameters	*`map-name`* — (Optional) Specify the name of a configured route map. A maximum of 140 characters.
Defaults	None

Command Mode EXEC

Usage Information None

Example

```
OS10# show route-map
route-map abc, permit, sequence 10
  Match clauses:
    ip address (access-lists): hello
    as-path abc
    community hello
    metric 2
    origin egp
    route-type external type-1
    tag 10
  Set clauses:
    metric-type type-1
    origin igp
    tag 100
```

Supported Releases 10.3.0E or later

12

Quality of service

Quality of service (QoS) reserves network resources for highly critical application traffic with precedence over less critical application traffic. QoS prioritizes different types of traffic and ensures quality of service.

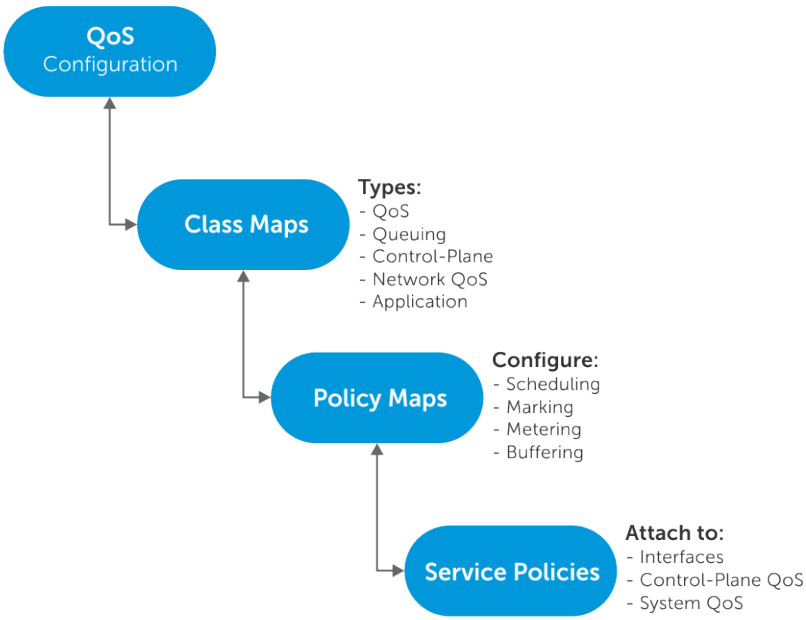
You can control the following traffic flow parameters: Delay, Bandwidth, Jitter, and Drop.

Different QoS features control the traffic flow parameters, as the traffic traverses a network device from ingress to egress interfaces.



Configure quality of service

Network traffic processes based on classification and policies that apply to the traffic.



Configuring QoS is a three-step process:

1 Create class-maps to classify the traffic flows. The following are the different types of class-maps:
- qos (default)—Classifies ingress data traffic.
- queuing —Classifies egress queues.
- control-plane—Classifies control-plane traffic.
- network-qos—Classifies traffic-class IDs for ingress buffer configurations.
- application —Classifies application-type traffic. The reserved policy-map **policy-iscsi** defines the actions for **class-iscsi** traffic.

2 Create policy-maps to define the policies for the classified traffic flows. The following are the different types of policy-maps:
- qos (default)—Defines the following actions on the traffic classified based on **qos** class-map:
  - Policing
  - Marking with a traffic class ID
  - Modifying packet fields such as CoS and DSCP
  - Enabling trust based classification
- queuing —Defines the following actions on the egress queues classified based on **queuing** class-map:
  - Shaping
  - Assigning bandwidth for queues
  - Assigning strict priority for queues
  - Buffering configuration for queues
  - WRED configuration on queues
- control-plane—Defines the policing of control queues for rate-limiting the **control-plane** traffic on CPU queues.
- network-qos—Defines the Ingress buffer configuration for selected traffic-classes matched based on **network-qos** class-map.
- application —Defines the following actions for the **application** classified traffic:
  - Modifying packet fields such as CoS and DSCP.
  - Marking traffic class IDs.

3 Apply the policy-maps to the port interface, system for all interfaces, or control-plane traffic as follows:
- Apply control-plane polices in Control-Plane mode.
- Apply QoS and network-QoS policies in the input direction on physical interfaces or in System-Qos mode.

- Apply queuing policies in the output direction on physical interfaces or in System-Qos mode.
- Apply a application type policy-map in System-Qos mode.

When you apply a policy on a system, the policy is effective on all the ports in the system. However, the interface-level policy takes precedence over the system-level policy.

Ingress traffic classification

Ingress traffic can either be data or control traffic.
By default, OS10 does not classify data traffic and assigns the default traffic class ID 0 to all data traffic.

OS10 implicitly classifies all control traffic such as STP, OSPF, ICMP, and so on, and forwards the traffic to control plane applications.

Data traffic classification

You can classify the data traffic based on ACL or trust.

ACL-based classification consumes significant amount of network processor resources. Trust-based classification classifies traffic in a pre-defined way without using network processor resources.

Trust based classification

OS10 supports classification based on the 802.1p CoS field (L2) or DSCP field (L3).
802.1p CoS trust map:

Trust the 802.1p CoS field to mark with a traffic-class ID and color for the CoS flow.

Table 60. Default 802.1p CoS trust map

CoS	Traffic class ID	Color
0	1	G
1	0	G
2	2	G
3	3	G
4	4	G
5	5	G
6	6	G
7	7	G

User–defined 802.1p CoS trust map

You can override the default mapping by creating a dot1p trust map. All the unspecified dot1p entries map to the default traffic class ID 0.

Configure user–defined 802.1p CoS trust map

1 Create a dot1p trust map.

```
OS10(config)# trust dot1p-map dot1p-trust-map
OS10(config-tmap-dot1p-map)#
```

2 Define the set of dot1p values mapped to traffic-class, the qos-group ID.

```
OS10(config-tmap-dot1p-map)# qos-group 3 dot1p 0-4
OS10(config-tmap-dot1p-map)# qos-group 5 dot1p 5-7
```

3 Verify the map entries.

```
OS10# show qos maps type trust-map-dot1p dot1p-trust-map

DOT1P Priority to Traffic-Class Map : dot1p-trust-map

Traffic-Class    DOT1P Priority

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

3                    0-4

5                    5-7
```

4 Apply the map on a specific interface or on system-qos, global level.

- Interface level

```
OS10(conf-if-eth1/1/1)# trust-map dot1p dot1p-trust-map
```

> **NOTE: In the interface level, the `no` version of the command returns the configuration to the system-qos level. If there is no configuration available at the system-qos level, the configuration returns to default mapping.**

- System-qos level

```
OS10(config-sys-qos)# trust-map dot1p dot1p-trust-map
```

Configure default CoS trust map

1 Create a default dot1p trust map.

```
OS10(config)# trust dot1p-map default
OS10(config-tmap-dot1p-map)#
```

2 Apply the map on a specific interface or on system-qos (global) level.

- Interface level

```
OS10(conf-if-eth1/1/1)# trust-map dot1p default
```

> **NOTE: In the interface level, the `no` version of the command returns the configuration to system-qos level. If there is no configuration available at the system-qos level, then the configuration returns to default mapping.**

- System-qos level

```
OS10(config-sys-qos)# trust-map dot1p default
```

DSCP trust map:

Assign a predefined and reserved trust classification in the policy map for the DSCP flow.

Table 61. Default DSCP trust map

DSCP values	TC id	Color
0-3	0	G
4-7	0	Y
8-11	1	G
12-15	1	Y
16-19	2	G
20-23	2	Y

DSCP values	TC id	Color
24-27	3	G
28-31	3	Y
32-35	4	G
36-39	4	Y
40-43	5	G
44-47	5	Y
48-51	6	G
52-55	6	Y
56-59	7	G
60-62	7	Y
63	7	R

User–defined DCSP trust map

Override the default mapping by creating a user-defined DSCP trust map. All the unspecified DSCP entries mapp to the default traffic class ID 0.

Configure user–defined DSCP trust map

1 Create a DSCP trust map.

```
OS10(config)# trust dscp-map dscp-trust-map
OS10(config-tmap-dscp-map)#
```

2 Define the set of dscp values mapped to traffic-class, the qos-group ID.

```
OS10(config-tmap-dscp-map)# qos-group 3 dscp 0-15
OS10(config-tmap-dscp-map)# qos-group 5 dscp 16-30
```

3 Verify the map entries.

```
OS10# show qos maps type trust-map-dscp dscp-trust-map

DSCP Priority to Traffic-Class Map  : dscp-trust-map

Traffic-Class    DSCP Priority

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

3                   0-15

5                   16-30
```

4 Apply the map on a specific interface or on system-qos global level.

- Interface level

  ```
  OS10(conf-if-eth1/1/1)# trust-map dscp dscp-trust-map
  ```

- System-qos level

  ```
  OS10(config-sys-qos)# trust-map dscp dscp-trust-map
  ```

Configure default DSCP trust map

1 Create a default DSCP trust map.

```
OS10(config)# trust dscp-map default
OS10(config-tmap-dscp-map)#
```

2 Apply the map on a specific interface or on system-qos global level.

- Interface level

  ```
  OS10(conf-if-eth1/1/1)# trust-map dscp default
  ```

- System-qos level

  ```
  OS10(config-sys-qos)# trust-map dscp default
  ```

ACL based classification

Classify the ingress traffic by matching the packet fields using ACL entries.

Classify the traffic flows based on QoS-specific fields or generic fields, using IP or MAC ACLs. Create a class-map template to match the fields.

OS10 allows matching *any* of the fields or *all* the fields based on the match type you configure in the class-map.

Use the access-group match filter to match MAC or IP ACLs. You can configure a maximum of four access-group filters in a class-map:

- 802.1p CoS
- VLAN ID (802.1.Q)
- DSCP + ECN
- IP precedence

OS10 supports configuring a range of or comma-separated values of match filters. When you apply the same match filter with new values, the system overwrites the previous values with the new values.

Configure ACL based classification

1 Create a class-map of type *qos*.

```
OS10(config)# class-map cmap
```

2 Define the fields to match, based on:

- 802.1p CoS

  ```
  OS10(config-cmap-qos)# match cos 0,4-7
  ```

- all the 802.1p CoS values excluding a few

  ```
  OS10(config-cmap-qos)# match not cos 3,4
  ```

- VLAN ID (range of or comma separated VLAN match is not supported)

  ```
  OS10(config-cmap-qos)# match vlan 100
  ```

- IP DSCP

  ```
  OS10(config-cmap-qos)# match ip dscp 3,5,20-30
  ```

- IP DSCP + ECN

  ```
  OS10(config-cmap-qos)# match ip dscp 3,5,20-30 ecn 2
  ```

- IP precedence

  ```
  OS10(config-cmap-qos)# match ip precedence 2
  ```

- IPv6 DSCP

  ```
  OS10(config-cmap-qos)# match ipv6 dscp 3,5,20-30
  ```

- IPv6 DSCP + ECN

  ```
  OS10(config-cmap-qos)# match ipv6 dscp 3,5,20-30 ecn 2
  ```

- IPv6 precedence

  ```
  OS10(config-cmap-qos)# match ipv6 precedence 2
  ```

- any IP (IPv4 or IPv6) precedence

  ```
  OS10(config-cmap-qos)# match ip-any precedence 2
  ```

- Pre-defined IP access-list

  ```
  OS10(config-cmap-qos)# match ip access-group name ip-acl-1
  ```

- Pre-defined IPv6 access-list

  ```
  OS10(config-cmap-qos)#match ipv6 access-group name ACLv6
  ```

- Pre-defined MAC access-list

  ```
  OS10(config-cmap-qos)# match mac access-group name mac-acl-1
  ```

3 Create a qos-type policy-map to refer the classes to.

```
OS10(config)# policy-map cos-policy
```

4 Refer the class-maps in the policy-map and define the required action for the flows.

```
OS10(config-pmap-qos)# class cmap
OS10(config-pmap-c-qos)# ?

OS10(config-pmap-qos)# class cmap
OS10(config-pmap-c-qos)#
  end     Exit to the exec Mode
  exit    Exit from current mode
  no      Negate a command or set its defaults
  police  Rate police input traffic
  set     Mark input traffic
  show    show configuration
  trust   Specify dynamic classification to trust[dscp/dot1p]
```

ACL based classification with trust

This section describes how to configure ACL based classification when you configure trust-based classification.

You can configure ACL based classification when trust-based classification is configured.

1 Create a user defined dscp or dot1p trust-map.

```
OS10(config)# trust dscp-map userdef-dscp

OS10(config-tmap-dscp-map)# qos-group 3 dscp 15

OS10(config-tmap-dscp-map)# qos-group 5 dscp 30
```

2 Apply user-defined trust map to an interface or in system QoS.

```
OS10(conf-if-eth1/1/1)# trust-map dscp userdef-dscp

or

OS10(config)# system qos

OS10(config-sys-qos)# trust-map dscp userdef-dscp
```

3 Create a class-map and attach it to a policy where fallback trust is configured.

```
OS10(config)# class-map c1
OS10(config-cmap-qos)# match cos 1
OS10(config-cmap-qos)# exit


OS10(config)# policy-map p1
OS10(config-pmap-qos)# class c1
OS10(config-pmap-c-qos)# set qos-group 1
```

4 Attach the policy map to an interface or in system QoS mode.

```
OS10(config)# interface ethernet 1/1/1

OS10(conf-if-eth1/1/1)# service-policy input type qos p1

or

OS10(config)# system qos

OS10(config-sys-qos)# service-policy input type qos p1
```

Control-plane policing

Control-plane policing (CoPP) increases security on the system by protecting the route processor from unnecessary traffic and giving priority to important control plane and management traffic. CoPP uses a dedicated control plane configuration through the QoS CLIs to set rate-limiting capabilities for control plane packets.

If the rate of control packets towards the CPU is higher than the packet rate that the CPU can handle, CoPP provides a method to selectively drop some of the control traffic so that the CPU can process high-priority control traffic. You can use CoPP to rate-limit traffic through each CPU port queue of the network processor (NPU).

CoPP applies policy actions on all control-plane traffic. The control-plane class map does not use any match criteria. To enforce rate-limiting or rate policing on control-plane traffic, create policy maps. You can use the `control-plane` command to attach the CoPP service policies directly to the control-plane.

Starting from release 10.4.2, the default rate limits change from 12 to 21 CPU queues and the protocols mapped to each CPU queue.

NOTE: When you upgrade from a previous release to release 10.4.2 and you have CoPP policy with rate limits configured in the previous release, the CoPP policies are automatically remapped based on the new CoPP protocol mappings to queues. For example:

- You have a CoPP policy configured for queue 5 in release 10.4.1, which is for ARP Request, ICMPv6-RS-NS, iSCSI snooping, and iSCSI-COS.
- After upgrade to release 10.4.2, the CoPP policy for queue 5 is remapped based on the new CoPP protocol mappings to queues as follows:
  - ARP Request is mapped to queue 6
  - ICMPv6-RS-NS is mapped to queue 5
  - iSCSI is mapped to queue 0

    The rate limit configuration in CoPP policy before upgrade is automatically remapped to queues 6, 5, and 0 respectively after upgrade.

For example, in release 10.4.1, the following policy configuration is applied on queue 5, which in 10.4.1 is mapped to ARP_REQ, ICMPV6_RS, ICMPV6_NS, and ISCSI protocols:

```
policy-map type control-plane test
!
class test
  set qos-group 5
  police cir 300 pir 300
```

After upgrade to release 10.4.2, the policy configuration appears as follows:

```
policy-map type control-plane test
!
class test_Remapped_0
  set qos-group 0
  police cir 300 pir 300
!
class test_Remapped_5
  set qos-group 5
  police cir 300 pir 300
!
class test_Remapped_6
  set qos-group 6
  police cir 300 pir 300
```

In release 10.4.2, ARP_REQ is mapped to queue 6, ICMPV6_RS and ICMPV6_NS are mapped to queue 5, and ISCSI is mapped to queue 0.

By default, CoPP traffic towards the CPU is classified into different queues as shown below.

Table 62. CoPP: Protocol mappings to queues - prior to release 10.4.2

Queue	Protocol
0	IPv6
1	—
2	IGMP
3	VLT, NDS
4	ICMPv6, ICMPv4
5	ARP Request, ICMPV6-RS-NS, ISCSI snooping, ISCSI-COS
6	ICMPv6-RA-NA, SSH, TELNET, TACACS, NTP, FTP
7	RSTP,PVST, MSTP,LACP
8	Dot1X,LLDP, FCOE-FPORT
9	BGPv4, OSPFv6
10	DHCPv6, DHCPv4, VRRP
11	OSPF Hello, OpenFlow

Table 63. CoPP: Protocol mappings to queues, and default rate limits and buffer sizes - from release 10.4.2 and later

Queue	Protocols	Minimum rate limit (in pps)	Maximum rate limit (in pps)	Minimum guaranteed buffer (in bytes)	Static shared limit (in bytes)
0	Unresolved, iSCSI, IPv6	600	600	1664	20800
1	SFlow	1000	1000	1664	20800
2	IGMP, MLD, PIM control	400	400	1664	48880
3	VLT, NDS	600	1000	1664	48880
4	IPv6 ICMP, IPv4 ICMP	500	500	1664	20800
5	ICMPv6 RS, RA, NS, NA	500	500	1664	48880
6	ARP request	500	1000	1664	48880
7	ARP response	500	1000	1664	48880
8	SSH, TELNET, NTP, FTP, TACACS	500	500	1664	20800
9	FCoE	600	600	1664	48880

Queue	Protocols	Minimum rate limit (in pps)	Maximum rate limit (in pps)	Minimum guaranteed buffer (in bytes)	Static shared limit (in bytes)
10	LACP	600	1000	1664	48880
11	STP, RSTP, MSTP	400	400	1664	48880
12	DOT1X, LLDP	500	500	1664	48880
13	IPv6 OSPF	600	1000	1664	48880
14	IPv4 OSPF	600	1000	1664	48880
15	BGP	600	1000	1664	48880
16	IPv4 DHCP, IPv6 DHCP	500	500	1664	48880
17	VRRP	600	1000	1664	48880
18	BFD	700	700	1664	48880
19	Remote CPS	700	1000	1664	48880
20	MCAST data	300	300	1664	20800
21	ACL logging	100	100	1664	20800
22	MCAST known data	300	300	1664	20800

For information about the current protocol to queue mapping and the rate-limit configured per queue, see show control-plane info.

Configure control-plane policing

Rate-limiting the protocol CPU queues requires configuring control-plane type QoS policies.

- Create QoS policies, class maps and policy maps, for the desired CPU-bound queue.
- Associate the QoS policy with a particular rate-limit.
- Assign the QoS service policy to control plane queues.

By default, the peak information rate (`pir`) and committed information rate (`cir`) values are in packets per second (pps) for control plane. CoPP for CPU queues converts the input rate from kilobits per second (kbps) to packets per second (pps), assuming 64 bytes is the average packet size, and applies that rate to the corresponding queue – One kbps is roughly equivalent to two pps.

1 Create a `control-plane` type class-map and configure a name for the class-map in CONFIGURATION mode.

```
class-map type control-plane class-map-name
```

2 Return to CONFIGURATION mode.

```
exit
```

3 Create an input policy-map to assign the QoS policy to the desired service queues in CONFIGURATION mode.

```
policy-map type control-plane policy-map-name
```

4 Associate a policy-map with a class-map in POLICY-MAP mode.

```
class class-name
```

5 Configure marking for a specific queue number in POLICY-MAP-CLASS-MAP mode, from 0 to 20.

```
set qos-group queue-number
```

6 Configure rate policing on incoming traffic in POLICY-MAP-CLASS-MAP mode.

```
police {cir committed-rate | pir peak-rate}
```

- `cir committed-rate`—Enter a committed rate value in pps, from 0 to 4000000.
- `pir peak rate` — Enter a peak-rate value in pps, from 0 to 40000000.

Create QoS policy for CoPP

```
OS10(config)# class-map type control-plane copp
OS10(conf-cmap-control-plane)# exit
OS10(config)# policy-map type control-plane copp1
OS10(conf-pmap-control-plane)# class copp
OS10(conf-pmap-c)# set qos-group 2
OS10(conf-pmap-c)# police cir 100 pir 100
```

View policy-map

```
OS10(conf-pmap-c)# do show policy-map
 Service-policy(control-plane) input: copp1
   Class-map (control-plane): copp
     set qos-group 2
     police cir 100 bc 100 pir 100 be 100
```

Assign service-policy

Rate controlling the traffic towards CPU requires configuring the **control-plane** type policy. To enable CoPP, apply the defined policy-map to CONTROL-PLANE mode.

1 Enter CONTROL-PLANE mode from CONFIGURATION mode.

```
control-plane
```

2 Define an`input` type service-policy and configure a name for the service policy in CONTROL-PLANE mode.

```
service-policy input service-policy-name
```

Assign control-plane service-policy

```
OS10(config)# control-plane
OS10(conf-control-plane)# service-policy input copp1
```

View control-plane service-policy

```
OS10(conf-control-plane)# do show qos control-plane
 Service-policy (input): copp1
```

View configuration

Use `show` commands to display the protocol traffic assigned to each control-plane queue and the current rate-limit applied to each queue. Use the `show` command output to verify the CoPP configuration.

View CoPP configuration

```
OS10# show qos control-plane
 Service-policy (input): pmap1
```

View CMAP1 configuration

```
OS10# show class-map type control-plane cmap1
 Class-map (control-plane): cmap1 (match-any)
```

View CoPP service-policy

```
OS10# show policy-map type control-plane
 Service-policy(control-plane) input: pmap1
   Class-map (control-plane): cmap1
```

```
    set qos-group 6
    police cir 200 bc 100 pir 200 be 100
```

View CoPP information

```
OS10# show control-plane info
Queue Min Rate Limit(in pps) Max Rate Limit(in pps)
Protocols
0     600                    600                    ISCSI UNKNOWN UNICAST
1     1000                   1000                   SFLOW
2     400                    400                    IGMP MLD PIM
3     600                    1000                   VLT NDS
4     500                    1000                   IPV6_ICMP IPV4_ICMP
5     500                    1000                   ICMPV6_RS ICMPV6_NS ICMPV6_RA ICMPV6_NA
6     500                    1000                   ARP_REQ SERVICEABILITY
7     500                    1000                   ARP_RESP
8     500                    500                    SSH TELNET TACACS NTP FTP
9     600                    600                    FCOE
10    600                    1000                   LACP
11    400                    400                    RSTP PVST MSTP
12    500                    500                    DOT1X LLDP
13    600                    1000                   IPV6_OSPF IPV4_OSPF
14    600                    1000                   OSPF_HELLO
15    600                    1000                   BGP
16    500                    500                    IPV6_DHCP IPV4_DHCP
17    600                    1000                   VRRP
18    700                    700                    BFD
19    700                    1000                   OPEN_FLOW REMOTE CPS
20    300                    300                    MCAST DATA
21    100                    100                    ACL LOGGING
22    300                    300                    MCAST KNOWN DATA
```

View CoPP statistics

```
OS10# show control-plane statistics
Queue               Packets             Bytes               Dropped Packets          Dropped Bytes
0                   26                  1768                      0                        0
1                   0                   0                         0                        0
2                   0                   0                         0                        0
3                   0                   0                         0                        0
4                   36                  3816                      0                        0
5                   36                  3096                      0                        0
6                   919                 58816                     0                        0
7                   67                  4288                      0                        0
8                   0                   0                         0                        0
9                   0                   0                         0                        0
10                  0                   0                         0                        0
11                  80662               5539376                   0                        0
12                  2779                462189                    0                        0
13                  0                   0                         0                        0
14                  1265                108790                    0                        0
15                  422                 36075                     0                        0
16                  0                   0                         0                        0
17                  0                   0                         0                        0
18                  0                   0                         0                        0
19                  0                   0                         0                        0
```

Egress traffic classification

Egress traffic is classified into different queues based on the traffic-class ID marked on the traffic flow.

Set the traffic class ID for a flow by enabling trust or by classifying ingress traffic and mark it with a traffic class ID using a policy map. By default, the value of traffic class ID for all the traffic is `0`.

The order of precedence for a qos-map is:

1 Interface-level map

2 System-qos-level map
3 Default map

Table 64. Default mapping of traffic class ID to queue

Traffic class ID	Queue ID
0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7

User–defined QoS map

You can override the default mapping by creating a QoS map.

Configure user–defined QoS map

1 Create a QoS map.
```
OS10(config)# qos-map traffic-class tc-q-map
```
2 Define the set of traffic class values mapped to a queue.
```
 OS10(config-qos-map)# queue 3 qos-group 0-3
```
3 Verify the map entries.
```
OS10# show qos maps type tc-queue

Traffic-Class to Queue Map: tc-q-map

Queue         Traffic-Class

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

3               0-3
```
4 Apply the map on a specific interface or on a system-QoS global level.
- Interface level
  ```
  OS10(conf-if-eth1/1/1)# qos-map traffic-class tc-q-map
  ```
- System-qos level
  ```
  OS10(config-sys-qos)# qos-map traffic-class tc-q-map
  ```

Choose all traffic classified for a queue

1 Create a queuing type class-map to match queue 5.
```
OS10(config)# class-map type queuing q5
```
2 Define the queue to match.
```
OS10(config-cmap-queuing)# match queue 5
```

Policing traffic

Use policing to limit the rate of ingress traffic flow. The flow can be all the ingress traffic on a port or a particular flow assigned with a traffic class ID.

In addition, use policing to color the traffic:

- When traffic arrives at a rate less than the committed rate, the color is green.
- When traffic propagates at an average rate greater than or equal to the committed rate and less than peak-rate, the color is yellow.
- When the traffic rate is above the configured peak-rate, the traffic drops to guarantee a bandwidth limit for an ingress traffic flow.

Peak rate is the maximum rate for traffic arriving or leaving an interface under normal traffic conditions. Peak burst size indicates the maximum size of unused peak bandwidth that is aggregated. This aggregated bandwidth enables brief durations of burst traffic that exceeds the peak rate.

Configure Interface rate policing

1 Create a QoS type empty class-map to match all the traffic.

```
OS10(config)# class-map cmap-all-traffic
```

2 Create a QoS type policy-map to define a policer.

```
OS10(config)# policy-map interface-policer
OS10(config-pmap-qos)# class cmap-all-traffic
OS10(config-pmap-c-qos)#police cir 4000 pir 6000
```

Configure flow rate policing

1 Create a QoS type class-map to match the traffic flow.

```
OS10(config)# class-map cmap-cos3
OS10(config-cmap-qos)# match cos 3
```

2 Create a QoS type policy-map to define a policer and assign a traffic class ID for the CoS flow.

```
OS10(config)# policy-map flow-policer
OS10(config-pmap-qos)# class cmap-cos3
OS10(config-pmap-c-qos)# set qos-group 3
OS10(config-pmap-c-qos)#police cir 4000 pir 6000
```

Mark Traffic

You can select a flow and mark it with a traffic class ID. Rraffic class IDs identify the traffic flow when the traffic reaches egress for queue scheduling.

Mark traffic

1 Create a QoS type class-map to match the traffic flow.

```
OS10(config)# class-map cmap-cos3
OS10(config-cmap-qos)# match cos 3
```

2 Create a QoS type policy-map to mark it with a traffic class ID and assign it to the CoS flow.

```
OS10(config)# policy-map cos3-TC3
OS10(config-pmap-qos)# class cmap-cos3
OS10(config-pmap-c-qos)# set qos-group 3
```

Color traffic

You can select a traffic flow and mark it with a color. Color the traffic flow based on:

- Metering. See Policing traffic.
- Default trust. See Trust-based classification.
- DSCP, ECN capable traffic (ECT), or non-ECT capable traffic.

Color traffic based on DSCP, ECT, or non-ECT

1 Create a QoS type class-map to match the traffic flow.

```
OS10(config)# class-map cmap-dscp-3-ect
OS10(config-cmap-qos)# match ip dscp 3 ecn 1
```

2 Create a QoS type policy-map to color the traffic flow.

```
OS10(config)# policy-map ect-color
OS10(config-pmap-qos)# class cmap-dscp-3-ect
OS10(config-pmap-c-qos)# set qos-group 3
OS10(config-pmap-c-qos)# set color yellow
```

Modify packet fields

You can modify the value of CoS or DSCP fields.

1 Create a QoS type class-map to match a traffic flow.

```
OS10(config)# class-map cmap-dscp-3
OS10(config-cmap-qos)# match ip dscp 3
```

2 Modify the policy-map to update the DSCP field.

```
OS10(config)# policy-map modify-dscp
OS10(config-pmap-qos)# class cmap-dscp-3
OS10(config-pmap-c-qos)# set qos-group 3
OS10(config-pmap-c-qos)# set dscp 10
```

Shaping traffic

You can shape the rate of egress traffic. When you enable rate shaping, the system buffers all traffic exceeding the specified rate until the buffer memory is exhausted. Rate shaping uses all buffers reserved for an interface or queue and shares buffer memory, until it reaches the configured threshold.

Configure traffic shaping

1 Enter the `queuing` type policy-map and configure a policy-map name in CONFIGURATION mode.

```
policy-map type queuing policy-map-name
```

2 Enter a class name to apply to the shape rate in POLICY-MAP-QUEUEING mode. A maximum of 32 characters.

```
class class-name
```

3 (Optional) Configure rate shaping on a specific queue by matching the corresponding qos-group in the class-map. If you do not configure the `match qos-group` command, rate shaping applies to all queues.

```
match qos-group queue-number
```

4 Enter a minimum and maximum shape rate value in POLICY-MAP-QUEUEING-CLASS mode.

```
shape {min {kbps | mbps}min-value} {max {kbps | mbps}max-value}
```

- 0 to 40000000—kilobits per second kilobits per second—kbps
- 0 to 40000 — megabits per second—mbps

Bandwidth allocation

You can allocate relative bandwidth to limit large flows and prioritize smaller flows. Allocate the relative amount of bandwidth to nonpriority queues when priorities queues are consuming maximum link bandwidth.

Schedule each egress queue of an interface per Weighted Deficit Round Robin (WDRR) or by strict-priority (SP), which are mutually exclusive. If the `bandwidth percent` command is present, you cannot configure the `priority` command.

In S5148F-ON, bandwidth weight is equally applied to UC and MC.

1 Create a `queuing` type class-map and configure a name for the class-map in CONFIGURATION mode.

```
class-map type queuing class-map-name
```

2 Apply the match criteria for the QoS group in CLASS-MAP mode.

```
qos-group queue-number
```

3 Return to CONFIGURATION mode.

```
exit
```

4 Create a `queuing` type policy-map and configure a policy-map name in CONFIGURATION mode.

```
policy-map type queuing policy-map-name
```

5 Configure a queuing class in POLICY-MAP mode.

```
class class-name
```

6 Assign a bandwidth percent, from 1 to 100 to nonpriority queues in POLICY-MAP-CLASS-MAP mode.

```
bandwidth percent value
```

Configure bandwidth allocation

```
OS10(config)# class-map type queuing solar
OS10(conf-cmap-queuing)# match qos-group 5
OS10(conf-cmap-queuing)# exit
OS10(config)# policy-map lunar
OS10(config)# policy-map type queuing lunar
OS10(conf-pmap-queuing)# class solar
OS10(conf-pmap-c-que)# bandwidth percent 80
```

View class-map

```
OS10(conf-cmap-queuing)# do show class-map
 Class-map (queuing): solar (match-any)
   Match: qos-group 5
```

View policy-map

```
OS10(conf-pmap-c-que)# do show policy-map
 Service-policy (queuing) output: solar
   Class-map (queuing): lunar
     bandwidth percent 80
```

Strict priority queuing

OS10 uses queues for egress QoS policy types. Enable priorities to dequeue all packets from the assigned queue before servicing any other queues. When you assign more than one queue strict priority, the highest number queue receives the highest priority. You can configure strict priority to any number of queues. By default, all queues schedule traffic per WDRR.

Use the `priority` command to assign the priority to a single unicast queue—this configuration supersedes the `bandwidth percent` configuration. A queue with priority enabled can starve other queues for the same egress interface.

Consider the following when enabling priority queueing in S5148F-ON:

- In a port, one H2 node and three H1 nodes are supported. The H1 node holds 8 unicast queues for data traffic, 8 unicast queues for control traffic, and 8 multicast queues for data traffic.
- The H1 nodes mapped to data traffic are scheduled with DWRR and weight of 50 each. The H1 node mapped to control traffic is scheduled with strict priority.
- The weights corresponding to each traffic class are applied at queue levels for both unicast and multicast queues.
- The bandwidth distribution might go to a minimum of 50, based on the traffic flow in a port. This is determined by the weight of a particular traffic class and traffic type.
- The bandwidth sharing based on ETS happens only between same type of queues.
- You can enable strict priority queuing only for the same type of traffic.

Create class-map

1 Create a class-map and configure a name for the class-map in CONFIGURATION mode.

```
class-map type queuing class-map-name
```

2 Configure a match criteria in CLASS-MAP mode.

```
match queue queue-id
```

Define a policy-map

1 Define a policy-map and create a policy-map name CONFIGURATION mode.

```
policy-map type queuing policy-map-name
```

2 Create a QoS class and configure a name for the policy-map in POLICY-MAP mode.

```
class class-map-name
```

3 Set the scheduler as strict priority in POLICY-MAP-CLASS-MAP mode.

```
priority
```

Apply policy-map

1 Apply the policy-map to the interface in INTERFACE mode or all interfaces in SYSTEM-QOS mode.

```
system qos
```

OR

```
interface ethernet node/slot/port[:subport]
```

2 Enter the output service-policy in SYSTEM-QOS mode or INTERFACE mode.

```
service-policy {output} type {queuing} policy-map-name
```

Enable strict priority on class-map

```
OS10(config)# class-map type queuing magnum
OS10(conf-cmap-queuing)# match queue 7
OS10(conf-cmap-queuing)# exit
OS10(config)# policy-map type queuing solar
OS10(conf-pmap-queuing)# class magnum
OS10(conf-pmap-c-que)# priority
OS10(conf-pmap-c-que)# exit
OS10(conf-pmap-queuing)# exit
OS10(config)# system qos
OS10(conf-sys-qos)# service-policy output solar
```

View QoS system

```
OS10(conf-sys-qos)# do show qos system
 Service-policy (output)(queuing): solar
```

Enable strict priority on interface

```
OS10(config)# interface ethernet 1/1/5
OS10(conf-if-eth1/1/5)# service-policy output type queuing solar
```

View policy-map

```
OS10(conf-if-eth1/1/5)# do show policy-map
 Service-policy(queuing) output: solar
   Class-map (queuing): magnum
     priority
```

Rate adjustment

QoS features such as policing and shaping do not include overhead fields such as Preamble, smart frame delimiter (SFD), inter-frame gap (IFG), and so on. For rate calculations, these feature only include the frame length between the destination MAC address (DMAC) and the CRC field.

You can optionally include the following overhead fields in rate calculations by enabling rate adjustment:

- Preamble—7 bytes
- Start frame delimiter—1 byte
- Destination MAC address—6 bytes
- Source MAC address—6 bytes
- Ethernet type/length—2 bytes

- Payload—variable
- Cyclic redundancy check—4 bytes
- Inter-frame gap—variable

The rate adjustment feature is disabled by default. To enable rate adjustment, use the `qos-rate-adjust value_of_rate_adjust` command. For example:

```
qos-rate-adjust 8
```

If you have configured WDRR and shaping on a particular queue, the queue can become congested. You should configure the QoS rate adjust value considering the overhead field size to avoid traffic drops on uncongested queues.

If you have multiple streams within a queue, you must find the overhead size for the different streams and the QoS rate adjust value should be the highest overhead size from among the various streams within that queue.

Consider the example where you have configured WDRR and shaping on a queue that has two different traffic streams, TS1 and TS2, that uses preamble, SFD, and IFG overhead fields:

- If the IFG in TS1 uses 16 bytes, QoS rate adjust value should be 24 (preamble + SFD requires 8 bytes and IFG 16 bytes).
- If the IFG in TS2 uses 12 bytes, QoS rate adjust value should be 20 (preamble + SFD requires 8 bytes and IFG 12 bytes).

In this case, the highest QoS rate adjust value between the two streams is 24 bytes. Hence, you must configure the QoS rate adjust value as 24.

NOTE: This feature is not supported on the S4200-ON Series and S5148F-ON platforms.

Buffer management

OS10 devices distribute the total available buffer resources into two buffer pools at ingress direction and three buffer pools at egress direction of all physical ports.

All ports in a system are allocated a certain amount of buffers from corresponding pools based on the configuration state of each priority-group or queue. The remaining buffers in the pool are shared across all similarly configured ports.

The following buffer pools are available:

- Ingress buffer pools:
  - Lossy pool (default)
  - Lossless pool (PFC)
- Egress buffer pools:
  - Lossy pool (default)
  - Lossless pool (PFC)
  - CPU pool (CPU control traffic)

For example, when all ports are allocated as reserved buffers from the lossy (default) pool, the remaining buffers in the lossy pool are shared across all ports, except the CPU port.

When you enable priority flow control (PFC) on the ports, all the PFC-enabled queues and priority-groups use the buffers from the lossless pool.

OS10 dedicates a separate buffer pool for CPU traffic. All default reserved buffers for the CPU port queues are from the CPU pool. The remaining buffers are shared across all CPU queues. You can modify the buffer settings of CPU queues.

You can configure the size of the CPU pool using the `control-plane-buffer-size` command.

OS10 allows configuration of buffers per priority-group and queue for each port.

Buffer-usage accounting happens for ingress packets on ingress pools and egress packets on egress pool. You can configure ingress-packet buffer accounting per priority-group and egress-packet buffer accounting per queue level.

Configure ingress buffer

In default ingress buffers, all traffic classes map to the default priority group. The buffers are reserved per default priority group ID 7. All buffers are part of the default pool and all ports share buffers from the default pool.

The reserved buffer size is 9360 bytes for the speed of 10G, 25G, 40G, 50G, and 100G. The supported speed varies for different platforms.

Table 65. Maximum buffer size

Platforms	Max buffer size
S4000	12 MB
S6010–ON, S4048–ON	16 MB
S4100-ON Series	12 MB
Z9100–ON	16 MB

The following lists the link-level flow control (LLFC) buffer settings for default priority group 7:

Table 66. Default setting for LLFC

Speed	10G	25G	40G	50G	100G
Default reserved buffer	9KB	9KB	18KB	18KB	36KB
Default Xon threshold	36KB	45KB	75KB	91KB	142KB
Default Xoff threshold	9KB	9KB	9KB	9KB	9KB
Default dynamic shared buffer threshold (alpha value)	9KB	9KB	9KB	9KB	9KB

NOTE: The supported speed varies for different platforms. After the reserved buffers are used, each LLFC starts consuming shared buffers from the lossless pool with the alpha value determining the threshold except for the S4200-ON series and S5148F-ON platform.

The following table lists the priority flow control (PFC) buffer settings per PFC priority group:

Table 67. Default settings for PFC

Speed	10G	25G	40G	50G	100G
Default reserved buffer for S4000, S4048–ON, S6010–ON	9KB	NA	9KB	NA	NA
Default reserved buffer for S41xx, Z9100–ON	9KB	9KB	18KB	18KB	36KB
Default Xoff threshold	36KB	45KB	75KB	91KB	142KB
Default Xon threshold	9KB	9KB	9KB	9KB	9KB
Default dynamic share buffer threshold (alpha value)	9KB	9KB	9KB	9KB	9KB

NOTE: The supported speed varies for different platforms. After the reserved buffers are used, each PFC starts consuming shared buffers from the lossless pool with the alpha value determining the threshold.

You can override the default priority group settings when you enable LLFC or PFC.

1 Create a network-qos type class-map to match the traffic classes. For LLFC, match all the traffic classes from 0 to 7. For PFC, match the required traffic class.

```
OS10(config)# class-map type network-qos tc
OS10 (config-cmap-nqos)# match qos-group 0-7
```

2 Create network-qos type policy-map to define the actions for traffic classes, such as a buffer configuration and threshold.

```
OS10(config)# policy-map type network-qos buffer
OS10(config-pmap-network-qos)# class tc
OS10 (config-pmap-c-nqos)# pause buffer-size 300 pause-threshold 200 resume-threshold 100
OS10 (config-pmap-c-nqos)# queue-limit thresh-mode dynamic 5
```

Configure egress buffer

All port queues are allocated with reserved buffers. When the reserved buffers are consumed, each queue starts using the shared buffers from the default pool.

The reserved buffer per queue is 1664 bytes for the speed of 10G, 25G, 40G, 50G, and 100G. The default dynamic shared buffer threshold is 8.

1 Create a queuing type class-map to match the queue.

```
OS10(config)# class-map type queuing q1
OS10(config-cmap-queuing)# match queue 1
```

2 Create a queuing type policy-map to define the actions for queues, such as a buffer configuration and threshold.

```
OS10(config)# policy-map type queuing q-buffer
OS10(config-pmap-queuing)# class q1
OS10(config-pmap-c-que)# queue-limit queue-len 200 thresh-mode dynamic 5
```

Deep Buffer mode

NOTE: This feature is supported only on the S4200-ON series.

OS10 provides the flexibility to configure the buffer mode based on your system requirements. The system memory contains a list of packet buffers and per packet information (PPI), which is used to enable statistics tagging, ingress shaping, PFC, and output logical interface stamping per multicast traffic.

You can configure Deep Buffer mode to manage switch buffer availability. The following lists the total buffer availability in different modes:

Table 68. Buffer availability in different modes

Platform	Without deep buffer	With deep buffer
S4200–ON series	4.63 GB	6.24 GB

Deep Buffer mode takes effect only after saving it in the startup configuration and reloading the switch.

NOTE: Disabling the Deep Buffer mode configuration during run-time is not supported.

Restrictions

Deep Buffer mode and network QoS configurations cannot coexist. Enable Deep Buffer mode only when the network QoS configurations; for example LLFC and PFC are disabled. To configure Deep Buffer mode, disable all network QoS related configurations.

Configure Deep Buffer mode

You must disable all the network QoS configurations; for example, PFC and LLFC, before configuring the Deep Buffer mode. Deep Buffer mode is disabled by default.

1 Enable Deep Buffer mode in CONFIGURATION mode.

```
OS10# configure terminal
OS10(config)# hardware deep-buffer-mode
```

> **NOTE: To disable Deep Buffer mode, use the `no` form of the command. Disabling Deep Buffer mode takes effect only after saving it in the startup configuration and reloading the switch.**

After you configure Deep Buffer mode, the system displays a warning stating that the configuration takes effect only after saving it in the startup configuration and reloading the switch.

```
% Warning: Deep buffer mode configuration will be applied only after a save and reload.
```

2 Save Deep Buffer mode in the startup configuration in CONFIGURATION mode.

```
OS10(config)# do write memory
```

3 Reload the switch in CONFIGURATION mode.

```
OS10(config)# do reload

Proceed to reboot the system? [confirm yes/no]:yes
```

View Deep Buffer mode status

To view Deep Buffer mode status, use the `show hardware deep-buffer-mode` command. The `show` command output displays the status of Deep Buffer mode in the current boot and the next boot.

The following is Deep Buffer mode status before enabling it, the default setting:

```
OS10# show hardware deep-buffer-mode
  Deep Buffer Mode Configuration Status
-----------------------------------------
    Current-boot Settings : Disabled

    Next-boot Settings    : Disabled
```

The following is Deep Buffer mode status after saving the configuration in the startup configuration:

```
OS10# show hardware deep-buffer-mode
  Deep Buffer Mode Configuration Status
-----------------------------------------
    Current-boot Settings : Disabled

    Next-boot Settings    : Enabled
```

The following is Deep Buffer mode status after the switch reloads:

```
OS10# show hardware deep-buffer-mode
  Deep Buffer Mode Configuration Status
-----------------------------------------
    Current-boot Settings : Enabled

    Next-boot Settings    : Enabled
```

Congestion avoidance

Congestion avoidance anticipates and takes necessary actions to avoid congestion. The following mechanisms avoid congestion:

- **Tail drop**—Packets are buffered at traffic queues. When the buffers are exhausted or reach the configured threshold, excess packets drop. By default, OS10 uses tail drop for congestion avoidance.
- **Random early detection (RED)**—In tail drop, different flows are not considered in buffer utilization. When multiple hosts start retransmission, tail drop causes TCP global re-synchronization. Instead of waiting for the queue to get filled up completely, RED starts dropping excess packets with a certain drop-probability when the average queue length exceeds the configured minimum threshold. The early drop ensures that only some of TCP sources slow down, which avoids global TCP re-synchronization.
- **Weighted random early detection (WRED)**—This allows different drop-probabilities and thresholds for each color — red, yellow, green — of traffic. You can configure the drop characteristics for three different flows by assigning the colors to the flow. Assign colors to a particular flow or traffic using various methods, such as ingress policing, qos input policy-maps, and so on.
- **Explicit congestion notification (ECN)**—This is an extension of WRED. Instead of dropping the packets when the average queue length crosses the minimum threshold values, ECN marks the Congestion Experienced (CE) bit of the ECN field in a packet as ECN-capable traffic (ECT).

1 Configure a WRED profile in CONFIGURATION mode.

```
OS10(config)# wred wred_prof_1
```

2 Configure WRED threshold parameters for different colors in WRED CONFIGURATION mode.

```
OS10(config-wred)# random-detect color yellow minimum-threshold 100 maximum-threshold 300
drop-probability 40
```

3 Configure the exponential weight value for the WRED profile in WRED CONFIGURATION mode.

```
OS10(config-wred)# random-detect weight 4
```

4 Enable ECN.

```
OS10(config-wred)# random-detect ecn
```

5 Enable WRED/ECN on a queue.

```
OS10(config)# class-map type queuing c1
OS10(config-cmap-queuing)# match queue 2
OS10(config-cmap-queuing)# exit
OS10(config)# policy-map type queuing p1
OS10(config-pmap-queuing)# class c1
OS10(config-pmap-c-que)# random-detect wred_prof_1
```

6 Enable WRED/ECN on a port.

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# random-detect wred_prof_1
```

7 Enable WRED/ECN on a service-pool.

```
OS10(config)# system qos
OS10(config-sys-qos)# random-detect pool 0  wred_prof_1
```

NOTE: On the S4200–ON Series platform, enable ECN globally only. Also, apply ECN configurations only at the queue level. You cannot configure ECN at the interface or service-pool levels. If you try to apply the ECN configuration at the interface or service-pool levels, the configuration is not accepted.

1 Configure a WRED profile in CONFIGURATION mode.

```
OS10(config)# wred wred_prof_1
```

2 Configure WRED threshold parameters for different colors in WRED CONFIGURATION mode.

```
OS10(config-wred)# random-detect color yellow minimum-threshold 100 maximum-threshold 300
drop-probability 40
```

3 Configure the exponential weight value for the WRED profile in WRED CONFIGURATION mode.

```
OS10(config-wred)# random-detect weight 4
```

4 Configure the ECN threshold parameters in WRED CONFIGURATION mode.

```
OS10(config-wred)#random-detect ecn minimum-threshold 100 maximum-threshold 300 drop-
probability 40
```

5 Exit WRED CONFIGURATION mode.

```
OS10(config-wred)#exit
```

6 Enter QOS POLICY-MAP mode and create a queuing policy type.

```
OS10(config)#policy-map type queuing pol-map-1
```

7 Create a QoS class for the queuing policy type.

```
OS10(config-pmap-queuing)#class default
```

8 Assign a WRED profile to the specified queue.

```
OS10(config-pmap-c-que)#random-detect prof1
```

9 Exit CLASS MAP and POLICY MAP modes.

```
OS10(config-pmap-c-que)#exit
OS10(config-pmap-queuing)#exit
```

10 Enter SYSTEM QOS mode.

```
OS10(config)#configure system-qos
```

11 Enable ECN globally.

```
OS10(config-sys-qos)#random-detect ecn
```

After you enable ECN globally, ECN marks the CE bit of the ECN field in a packet as ECT.

In the S4200–ON Series platform, configure separate thresholds for ECN capable traffic (ECT). If you enable ECN, ECT is marked based on the configured ECN threshold and non-ECT drops based on the WRED thresholds.

Storm control

Traffic storms created by packet flooding or other reasons may degrade the performance of the network. The storm control feature allows you to control unknown unicast, multicast, and broadcast traffic on L2 and L3 physical interfaces.

In the storm control unknown unicast configuration, both the unknown unicast and unknown multicast traffic are rate-limited.

OS10 devices monitor the current level of the traffic rate at fixed intervals, compares the traffic rate with the configured levels, and drops excess traffic.

By default, storm control is disabled on all interfaces. Enable storm control using the `storm-control { broadcast | multicast | unknown-unicast }` *`rate-in-pps`* command in INTERFACE mode.

NOTE: On the S5148F-ON platform, there is a 2% of deviation in storm control configuration.

- Enable broadcast storm control with a rate of 1000 packets per second (pps) on Ethernet 1/1/1.

```
OS10(conf-if-eth1/1/1)# storm-control broadcast 1000
```

RoCE for faster access and lossless connectivity

Remote Direct Memory Access (RDMA) enables memory transfers between two computers in a network without involving the CPU of either computer.

RDMA networks provide high bandwidth and low latency without appreciable CPU overhead for improved application performance, storage and data center utilization, and simplified network management. RDMA was traditionally supported only in an InfiniBand environment. Currently, RDMA over Converged Ethernet (RoCE) is also implemented in data centers that use Ethernet or a mixed-protocol environment.

OS10 devices support RoCE v1 and RoCE v2 protocols.

- RoCE v1 – An Ethernet layer protocol that allows for communication between two hosts that are in the same Ethernet broadcast domain.
- RoCE v2 – An Internet layer protocol that allows RoCE v2 packets to be routed, called Routable RoCE (RRoCE).

To enable RRoCE, configure the QoS service policy on the switch in ingress and egress directions on all the interfaces. For more information about this configuration, see Configure RoCE on the switch.

Configure RoCE on the switch

The following example describes the steps to configure RoCE on the switch. This configuration example uses priority 3 for RoCE.

1 Enter CONFIGURATION mode.

```
OS10# configure terminal
OS10 (config)#
```

2 Enable the Data Center Bridging Exchange protocol (DCBX).

```
OS10 (config)# dcbx enable
```

3 Create a VLAN. In this example, VLAN 55 switchs the RoCE traffic. You can configure any value from 1 to 4093.

```
OS10 (config)# interface vlan 55
```

4 Create a network-qos type class-map for priority flow control (PFC).

```
OS10 (config)# class-map type network-qos pfcdot1p3
OS10 OS10(config-cmap-nqos)# match qos-group 3
```

5 Create queuing-type class-maps for enhanced transmission selection (ETS).

```
OS10 (config)# class-map type queuing Q0
OS10 (config)# match queue 0
OS10 (config)# class-map type queuing Q3
OS10 (config)# match queue 3
```

6 Create a QoS map for ETS.

```
OS10 (config)# qos-map traffic-class 2Q
OS10(config-qos-map)# queue 0 qos-group 0-2, 4-7
OS10(config-qos-map)# queue 3 qos-group 3
```

7 Create a policy-map for PFC.

```
OS10 (config)# policy-map type network-qos pfcdot1p3
OS10(config-pmap-network-qos)# class pfcdot1p3
OS10(config-pmap-c-nqos)# pause
```

8 Create an egress policy-map.

```
OS10 (config)# policy-map type queuing 2Q
OS10 (config-pmap-queuing)# class Q0
OS10 (config-pmap-c-que)# bandwidth percent 30
OS10 (config-pmap-c-que)# exit
OS10 (config-pmap-queuing)# class Q3
OS10 (config-pmap-c-que)# bandwidth percent 70
OS10 (config-pmap-c-que)#
```

9 Apply the dot1p trust globally or at the interface level. In this example, the dot1p trust is applied globally.

```
OS10 (config)# system qos
OS10 (config)# trust-map dot1p default
```

10 Perform the following configurations on all switch interfaces where you want to support RoCE:

a Enter INTERFACE mode and enter the `no shutdown` command.

```
OS10# configure terminal
OS10 (config)# interface ethernet 1/1/1
OS10 (conf-if-eth1/1/1)# no shutdown
```

b Change the switch port mode to Trunk mode.

```
OS10 (conf-if-eth1/1/1)# switchport mode trunk
```

c Specify the allowed VLANs on the trunk port.

```
OS10 (conf-if-eth1/1/1)# switchport trunk allowed vlan 55
```

d Apply the network-qos type policy-map to the interface.

```
OS10 (conf-if-eth1/1/1)# service-policy input type network-qos pfcdot1p3
```

e Apply the queuing policy to egress traffic on the interface.

```
OS10 (conf-if-eth1/1/1)# service-policy output type queuing 2Q
```

f Enable ETS on the interface.

```
OS10 (conf-if-eth1/1/1)# ets mode on
```

g Apply the qos-map for ETS configurations on the interface.

```
OS10 (conf-if-eth1/1/1)# qos-map traffic-class 2Q
```

h Enable PFC on the interface.

```
OS10 (conf-if-eth1/1/1)# priority-flow-control mode on
```

Configure RoCE on the interfaces

The following example describes the steps that you need to perform to configure RoCE on the all the interfaces that the switch uses. This configuration example uses priority 3 for RoCE.

1 Enter in to the INTERFACE mode and enter the `no shutdown` command.

```
OS10# configure terminal
OS10 (config)# interface ethernet 1/1/1
OS10 (conf-if-eth1/1/1)# no shutdown
```

2 Change the switch port mode to trunk mode.

```
OS10 (conf-if-eth1/1/1)# switchport mode trunk
```

3 Change the access VLAN management.

```
OOS10 (conf-if-eth1/1/1)# switchport access vlan 1
```

4 Specify the allowed VLANs on the trunk port.

```
OS10 (conf-if-eth1/1/1)# switchport trunk allowed vlan 55
```

5 Apply the policy-map to the interface.

```
OS10 (conf-if-eth1/1/1)# service-policy input type network-qos pfcdot1p3
```

6 Apply the queuing policy to egress traffic on the interface.

```
OS10 (conf-if-eth1/1/1)# service-policy output type queuing 2Q
```

7 Enable enhanced transmission selection (ETS) on the interface.

```
OS10 (conf-if-eth1/1/1)# ets mode on
```

8 Apply the service policy with dot1p trust and ETS configurations on the interface.

```
OS10 (conf-if-eth1/1/1)# qos-map traffic-class 2Q
```

9 Enable PFC on the interface.

```
OS10 (conf-if-eth1/1/1)# priority-flow-control mode on
```

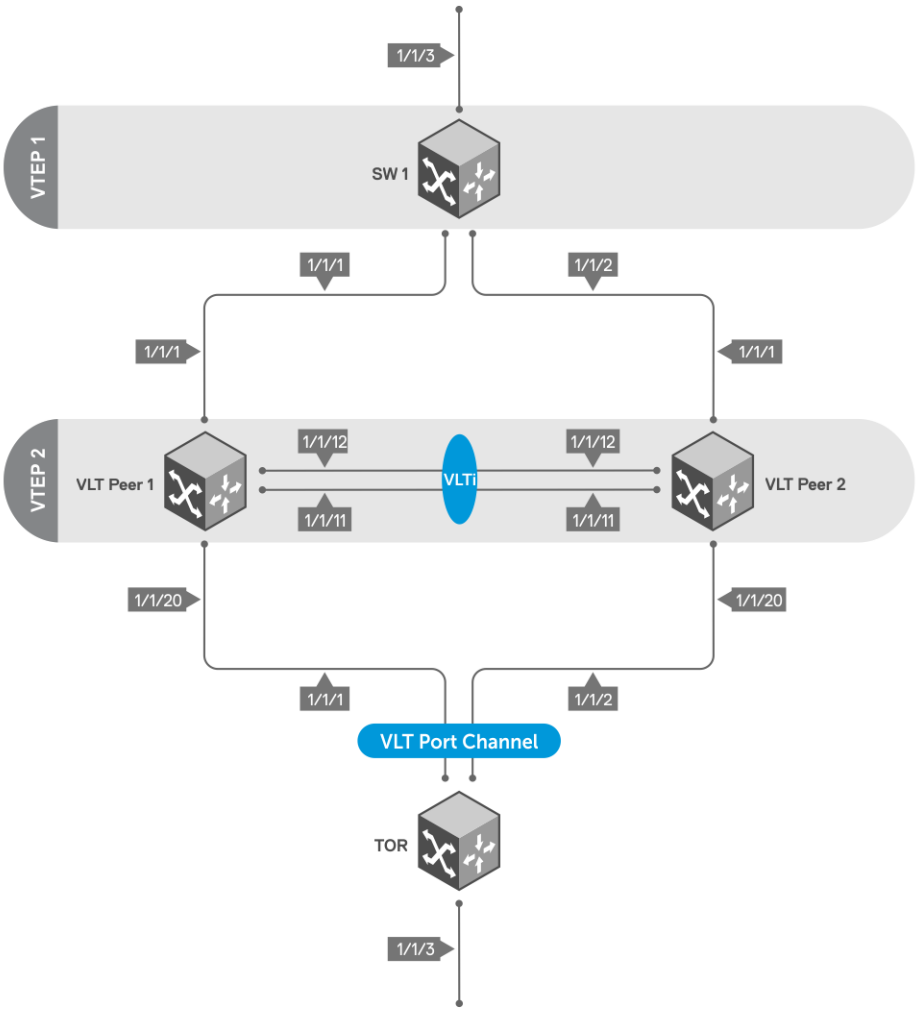
RoCE for VXLAN over VLT

OS10 supports RoCE for VXLAN in a VLT setup. Configuring RoCE with VXLAN is similar to configuring RoCE without VXLAN. When you configure VXLAN and span that across a VLT topology, apply the configuration on all interfaces across the VLT topology where you want to support RoCE.

For more information about how to configure RoCE, see the Configure RoCE on the switch section.

Sample configuration of RoCE for VXLAN over VLT

The following describes a topology where RoCE is enabled with VXLAN over VLT. SW1 is configured as VTEP1 and is the upstream switch that connects to the outer network. VLT peer 1 and VLT peer 2 form a VLT topology which is also configured as VTEP 2. A top-of-rack (ToR) switch is connected to the VLT peers through a VLT port channel. The ToR is the downstream switch for end devices.



The following examples show each device in this network and their respective configuration.

SW1 configuration

VXLAN configuration — SW1

```
OS10# configure terminal
OS10(config)# interface vlan 3000
OS10(conf-if-vl-3000)# exit
OS10(config)# interface vlan 200
OS10(conf-if-vl-200)# exit
OS10(config)# interface loopback 1
OS10(conf-if-lo-1)# ip address 1.1.1.1/32
OS10(conf-if-lo-1)# exit
OS10(config)# router ospf 1
OS10(config-router-ospf-1)# router-id 8.8.8.8
OS10(config-router-ospf-1)# exit
OS10(config)# interface vlan 3000
OS10(conf-if-vl-3000)# ip ospf 1 area 0
OS10(conf-if-vl-3000)# exit
OS10(config)# interface loopback 1
OS10(conf-if-lo-1)# ip ospf 1 area 0
OS10(conf-if-lo-1)# exit
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# switchport mode trunk
OS10(conf-if-eth1/1/1)# switchport trunk allowed vlan 3000
```

```
OS10(conf-if-eth1/1/1)# exit
OS10(config)# interface ethernet 1/1/2
OS10(conf-if-eth1/1/2)# switchport mode trunk
OS10(conf-if-eth1/1/2)# switchport trunk allowed vlan 3000
OS10(conf-if-eth1/1/2)# exit
OS10(config)# configure terminal
OS10(config)# nve
OS10(conf-nve)# source-interface loopback 1
OS10(conf-nve)# exit
OS10(config)# virtual-network 5
OS10(conf-vn-5)# vxlan-vni 1000
OS10(conf-vn-vxlan-vni)# remote-vtep 2.2.2.2
OS10(conf-vn-vxlan-vni-remote-vtep)# exit
OS10(conf-vn-vxlan-vni)# exit
OS10(conf-vn-5)# exit
OS10(config)# interface vlan 200
OS10(conf-if-vl-200)# virtual-network 5
OS10(conf-if-vl-200)# end
OS10#
OS10# configure terminal
OS10(config)# interface ethernet 1/1/3
OS10(conf-if-eth1/1/3)# switchport mode trunk
OS10(conf-if-eth1/1/3)# switchport trunk allowed vlan 200
OS10(conf-if-eth1/1/3)# end
```

PFC configuration — SW1

```
OS10# configure terminal
OS10(config)# trust dot1p-map t1
OS10(config-tmap-dot1p-map)# qos-group 0 dot1p 0
OS10(config-tmap-dot1p-map)# qos-group 1 dot1p 1
OS10(config-tmap-dot1p-map)# qos-group 2 dot1p 2
OS10(config-tmap-dot1p-map)# qos-group 3 dot1p 3
OS10(config-tmap-dot1p-map)# qos-group 4 dot1p 4
OS10(config-tmap-dot1p-map)# qos-group 5 dot1p 5
OS10(config-tmap-dot1p-map)# qos-group 6 dot1p 6
OS10(config-tmap-dot1p-map)# qos-group 7 dot1p 7
OS10(config-tmap-dot1p-map)# end
OS10# configure terminal
OS10(config)# class-map type network-qos c5
OS10(config-cmap-nqos)# match qos-group 5
OS10(config-cmap-nqos)# exit
OS10(config)# policy-map type network-qos p5
OS10(config-pmap-network-qos)# class c5
OS10(config-pmap-c-nqos)# pause
OS10(config-pmap-c-nqos)# pfc-cos 5
OS10(config-pmap-c-nqos)# end
OS10# configure terminal
OS10(config)# interface range ethernet 1/1/3,1/1/1,1/1/2
OS10(conf-range-eth1/1/3,1/1/1,1/1/2)# flowcontrol receive off
OS10(conf-range-eth1/1/3,1/1/1,1/1/2)# priority-flow-control mode on
OS10(conf-range-eth1/1/3,1/1/1,1/1/2)# ets mode on
OS10(conf-range-eth1/1/3,1/1/1,1/1/2)# service-policy input type network-qos p5
OS10(conf-range-eth1/1/3,1/1/1,1/1/2)# trust-map dot1p t1
```

LLFC configuration — SW1

Instead of PFC, you can configure LLFC as follows:

```
OS10(config)# configure terminal
OS10(config)# class-map type network-qos llfc
OS10(config-cmap-nqos)# match qos-group 0-7
OS10(config-cmap-nqos)# exit
OS10(config)# policy-map type network-qos llfc
OS10(config-pmap-network-qos)# class llfc
OS10(config-pmap-c-nqos)# pause buffer-size 100 pause-threshold 50 resume-threshold 10
OS10(config-pmap-c-nqos)# end
OS10#
OS10# configure terminal
```

```
OS10(config)# interface range ethernet 1/1/1,1/1/2,1/1/3
OS10(conf-range-eth1/1/1,1/1/2,1/1/3)# flowcontrol transmit on
OS10(conf-range-eth1/1/1,1/1/2,1/1/3)# flowcontrol receive on
OS10(conf-range-eth1/1/1,1/1/2,1/1/3)# service-policy input type network-qos llfc
OS10(conf-range-eth1/1/1,1/1/2,1/1/3)# end
```

WRED and ECN configuration — SW1

```
OS10# configure terminal
OS10(config)# wred w1
OS10(config-wred)# random-detect ecn
OS10(config-wred)# random-detect color green minimum-threshold 100 maximum-threshold 500 drop-
probability 100
OS10(config-wred)# random-detect color yellow minimum-threshold 100 maximum-threshold 500 drop-
probability 100
OS10(config-wred)# random-detect color red minimum-threshold 100 maximum-threshold 500 drop-
probability 100
OS10(config-wred)# exit
OS10(config)# class-map type queuing cq
OS10(config-cmap-queuing)# match queue 5
OS10(config-cmap-queuing)# exit
OS10(config)# policy-map type queuing pq
OS10(config-pmap-queuing)# class cq
OS10(config-pmap-c-que)# random-detect w1
OS10(config-pmap-c-que)# end
OS10# configure terminal
OS10(config)# interface range ethernet 1/1/3,1/1/1,1/1/2
OS10(conf-range-eth1/1/3,1/1/1,1/1/2)# flowcontrol receive off
OS10(conf-range-eth1/1/3,1/1/1,1/1/2)# priority-flow-control mode on
OS10(conf-range-eth1/1/3,1/1/1,1/1/2)# ets mode on
OS10(conf-range-eth1/1/3,1/1/1,1/1/2)# service-policy input type network-qos p5
OS10(conf-range-eth1/1/3,1/1/1,1/1/2)# service-policy output type queuing pq
OS10(conf-range-eth1/1/3,1/1/1,1/1/2)# trust-map dot1p t1
OS10(conf-range-eth1/1/3,1/1/1,1/1/2)# end
```

Enable DCBx — SW1

```
OS10# configure terminal
OS10(config)# dcbx enable
```

Configuration on VLT peer 1

VLT configuration — VLT peer 1

```
OS10# configure terminal
OS10(config)# interface range ethernet 1/1/12,1/1/11
OS10(conf-range-eth1/1/12,1/1/11)# no switchport mode
OS10(conf-range-eth1/1/12,1/1/11)# no switchport
OS10(conf-range-eth1/1/12,1/1/11)# no negotiation
OS10(conf-range-eth1/1/12,1/1/11)# exit
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# discovery-interface ethernet 1/1/12
OS10(conf-vlt-1)# discovery-interface ethernet 1/1/11
OS10(conf-vlt-1)# vlt-mac aa:bb:cc:dd:ee:ff
OS10(conf-vlt-1)# end
OS10#
OS10# configure terminal
OS10(config)# interface port-channel 2
OS10(conf-if-po-2)# vlt-port-channel 20
OS10(conf-if-po-2)# no shutdown
OS10(conf-if-po-2)# exit
OS10(config)# interface range ethernet 1/1/20
OS10(conf-range-eth1/1/20)# channel-group 2 mode active
OS10(conf-range-eth1/1/20)# exit
```

VXLAN configuration — VLT peer 1

```
OS10(config)# configure terminal
OS10(config)# interface vlan 3000
```

```
OS10(conf-if-vl-3000)# ip address 5.5.5.2/24
OS10(conf-if-vl-3000)# exit
OS10(config)# interface vlan 200
OS10(conf-if-vl-200)# exit
OS10(config)# interface loopback1
OS10(conf-if-lo-1)# no shutdown
OS10(conf-if-lo-1)# ip address 2.2.2.2/11
OS10(conf-if-lo-1)# exit
OS10(config)# router ospf 1
OS10(config-router-ospf-1)# router-id 9.9.9.9
OS10(config-router-ospf-1)# exit
OS10(config)# interface loopback 1
OS10(conf-if-lo-1)# ip ospf 1 area 0
OS10(conf-if-lo-1)#
OS10(conf-if-lo-1)# configure terminal
OS10(config)# interface vlan 3000
OS10(conf-if-vl-3000)# ip ospf 1 area 0
OS10(conf-if-vl-3000)# end
OS10# configure terminal
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# switchport mode trunk
OS10(conf-if-eth1/1/1)# switchport trunk allowed vlan 3000
OS10(conf-if-eth1/1/1)# exit
OS10(config)# nve
OS10(conf-nve)# source-interface loopback 1
OS10(conf-nve)# exit
OS10(config)# virtual-network 5
OS10(conf-vn-5)# vxlan-vni 1200
OS10(conf-vn-vxlan-vni)# remote-vtep 1.1.1.1
OS10(conf-vn-vxlan-vni-remote-vtep)# exit
OS10(conf-vn-vxlan-vni)# exit
OS10(conf-vn-5)# exit
OS10(config)# interface vlan 200
OS10(conf-if-vl-200)# virtual-network 5
OS10(conf-if-vl-200)# end
OS10#
OS10# configure terminal
OS10(config)# interface port-channel 2
OS10(conf-if-po-2)# switchport mode trunk
OS10(conf-if-po-2)# switchport trunk allowed vlan 200
OS10(conf-if-po-2)# end
```

PFC configuration — VLT peer 1

```
OS10# configure terminal
OS10(config)# trust dot1p-map t1
OS10(config-tmap-dot1p-map)# qos-group 0 dot1p 0
OS10(config-tmap-dot1p-map)# qos-group 1 dot1p 1
OS10(config-tmap-dot1p-map)# qos-group 2 dot1p 2
OS10(config-tmap-dot1p-map)# qos-group 3 dot1p 3
OS10(config-tmap-dot1p-map)# qos-group 4 dot1p 4
OS10(config-tmap-dot1p-map)# qos-group 5 dot1p 5
OS10(config-tmap-dot1p-map)# qos-group 6 dot1p 6
OS10(config-tmap-dot1p-map)# qos-group 7 dot1p 7
OS10(config-tmap-dot1p-map)# end
OS10# configure terminal
OS10(config)# class-map type network-qos c5
OS10(config-cmap-nqos)# match qos-group 5
OS10(config-cmap-nqos)# exit
OS10(config)# policy-map type network-qos p5
OS10(config-pmap-network-qos)# class c5
OS10(config-pmap-c-nqos)# pause
OS10(config-pmap-c-nqos)# pfc-cos 5
OS10(config-pmap-c-nqos)# end
OS10# configure terminal
OS10(config)# interface range ethernet 1/1/1,1/1/20,1/1/31,1/1/32
OS10(conf-range-eth1/1/1,1/1/20,1/1/31,1/1/32)# flowcontrol receive off
OS10(conf-range-eth1/1/1,1/1/20,1/1/31,1/1/32)# priority-flow-control mode on
OS10(conf-range-eth1/1/1,1/1/20,1/1/31,1/1/32)# ets mode on
```

```
OS10(conf-range-eth1/1/1,1/1/20,1/1/31,1/1/32)# service-policy input type network-qos p5
OS10(conf-range-eth1/1/1,1/1/20,1/1/31,1/1/32)# trust-map dot1p t1
OS10(conf-range-eth1/1/1,1/1/20,1/1/31,1/1/32)# end
```

LLFC configuration — VLT peer 1

Instead of PFC, you can configure LLFC as follows:

```
OS10# configure terminal
OS10(config)# class-map type network-qos llfc
OS10(config-cmap-nqos)# match qos-group 0-7
OS10(config-cmap-nqos)# exit
OS10(config)# policy-map type network-qos llfc
OS10(config-pmap-network-qos)# class llfc
OS10(config-pmap-c-nqos)# pause buffer-size 120 pause-threshold 50 resume-threshold 12
OS10(config-pmap-c-nqos)# end
OS10# configure terminal
OS10(config)# interface range ethernet 1/1/1,1/1/20,1/1/31,1/1/32
OS10(conf-range-eth1/1/1,1/1/20,1/1/31,1/1/32)# flowcontrol transmit on
OS10(conf-range-eth1/1/1,1/1/20,1/1/31,1/1/32)# flowcontrol receive on
OS10(conf-range-eth1/1/1,1/1/20,1/1/31,1/1/32)# service-policy input type network-qos llfc
OS10(conf-range-eth1/1/1,1/1/20,1/1/31,1/1/32)# end
```

WRED/ECN configuration — VLT peer 1

```
OS10# configure terminal
OS10(config)# wred w1
OS10(config-wred)# random-detect ecn
OS10(config-wred)# random-detect color green minimum-threshold 120 maximum-threshold 500 drop-
probability 100
OS10(config-wred)# random-detect color yellow minimum-threshold 120 maximum-threshold 500 drop-
probability 100
OS10(config-wred)# random-detect color red minimum-threshold 120 maximum-threshold 500 drop-
probability 100
OS10(config-wred)# exit
OS10(config)# class-map type queuing cq
OS10(config-cmap-queuing)# match queue 5
OS10(config-cmap-queuing)# exit
OS10(config)# policy-map type queuing pq
OS10(config-pmap-queuing)# class cq
OS10(config-pmap-c-que)# random-detect w1
OS10(config-pmap-c-que)# end
OS10# configure terminal
OS10(config)# interface range ethernet 1/1/1,1/1/20,1/1/12,1/1/11
OS10(conf-range-eth1/1/1,1/1/20,1/1/12,1/1/11)# service-policy input type network-qos p5
OS10(conf-range-eth1/1/1,1/1/20,1/1/12,1/1/11)# service-policy output type queuing pq
OS10(conf-range-eth1/1/1,1/1/20,1/1/12,1/1/11)# trust-map dot1p t1
OS10(conf-range-eth1/1/1,1/1/20,1/1/12,1/1/11)# flowcontrol receive off
OS10(conf-range-eth1/1/1,1/1/20,1/1/12,1/1/11)# priority-flow-control mode on
OS10(conf-range-eth1/1/1,1/1/20,1/1/12,1/1/11)# ets mode on
OS10(conf-range-eth1/1/1,1/1/20,1/1/12,1/1/11)# end
```

Enable DCBx — VLT peer 1

```
OS10# configure terminal
OS10(config)# dcbx enable
```

Configuration on VLT peer 2

VLT configuration — VLT peer 2

```
OS10# configure terminal
OS10(config)# interface range ethernet 1/1/11,1/1/12
OS10(conf-range-eth1/1/11,1/1/12)# no switchport mode
OS10(conf-range-eth1/1/11,1/1/12)# no switchport
OS10(conf-range-eth1/1/11,1/1/12)# no negotiation
OS10(conf-range-eth1/1/11,1/1/12)# exit
OS10(config)# vlt-domain 1
```

```
OS10(conf-vlt-1)# discovery-interface ethernet 1/1/11
OS10(conf-vlt-1)# discovery-interface ethernet 1/1/12
OS10(conf-vlt-1)# vlt-mac aa:bb:cc:dd:ee:ff
OS10(conf-vlt-1)# end
OS10#
OS10# configure terminal
OS10(config)# interface port-channel 2
OS10(conf-if-po-2)# vlt-port-channel 20
OS10(conf-if-po-2)# no shutdown
OS10(conf-if-po-2)# exit
```

VXLAN configuration — VLT peer 2

```
OS10(config)# configure terminal
OS10(config)# interface vlan 3000
OS10(conf-if-vl-3000)# ip address 5.5.5.3/24
OS10(conf-if-vl-3000)# exit
OS10(config)# interface vlan 200
OS10(conf-if-vl-200)# exit
OS10(config)# interface loopback 1
OS10(conf-if-lo-1)# no shutdown
OS10(conf-if-lo-1)# ip address 2.2.2.2/32
OS10(conf-if-lo-1)# exit
OS10(config)# router ospf 1
OS10(config-router-ospf-1)# router-id 10.10.10.10
OS10(config-router-ospf-1)# exit
OS10(config)# interface loopback 1
OS10(conf-if-lo-1)# ip ospf 1 area 0
OS10(conf-if-lo-1)# configure terminal
OS10(config)# interface vlan 3000
OS10(conf-if-vl-3000)# ip ospf 1 area 0
OS10(conf-if-vl-3000)# end
OS10# configure terminal
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# switchport mode trunk
OS10(conf-if-eth1/1/1)# switchport trunk allowed vlan 3000
OS10(conf-if-eth1/1/1)# exit
OS10(config)# nve
OS10(conf-nve)# source-interface loopback 1
OS10(conf-nve)# exit
OS10(config)# virtual-network 5
OS10(conf-vn-5)# vxlan-vni 1000
OS10(conf-vn-vxlan-vni)# remote-vtep 1.1.1.1
OS10(conf-vn-vxlan-vni-remote-vtep)# exit
OS10(conf-vn-vxlan-vni)# exit
OS10(conf-vn-5)# exit
OS10(config)# interface vlan 200
OS10(conf-if-vl-200)# virtual-network 5
OS10(conf-if-vl-200)# end
OS10#
OS10# configure terminal
OS10(config)# interface port-channel 2
OS10(conf-if-po-2)# switchport mode trunk
OS10(conf-if-po-2)# switchport trunk allowed vlan 200
OS10(conf-if-po-2)# end
```

PFC configuration — VLT peer 2

```
OS10# configure terminal
OS10(config)# trust dot1p-map t1
OS10(config-tmap-dot1p-map)# qos-group 0 dot1p 0
OS10(config-tmap-dot1p-map)# qos-group 1 dot1p 1
OS10(config-tmap-dot1p-map)# qos-group 2 dot1p 2
OS10(config-tmap-dot1p-map)# qos-group 3 dot1p 3
OS10(config-tmap-dot1p-map)# qos-group 4 dot1p 4
OS10(config-tmap-dot1p-map)# qos-group 5 dot1p 5
OS10(config-tmap-dot1p-map)# qos-group 6 dot1p 6
OS10(config-tmap-dot1p-map)# qos-group 7 dot1p 7
OS10(config-tmap-dot1p-map)# end
```

```
OS10# configure terminal
OS10(config)# class-map type network-qos c5
OS10(config-cmap-nqos)# match qos-group 5
OS10(config-cmap-nqos)# exit
OS10(config)# policy-map type network-qos p5
OS10(config-pmap-network-qos)# class c5
OS10(config-pmap-c-nqos)# pause
OS10(config-pmap-c-nqos)# pfc-cos 5
OS10(config-pmap-c-nqos)# end
OS10# configure terminal
OS10(config)# interface range ethernet 1/1/1,1/1/20,1/1/11,1/1/12
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# flowcontrol receive off
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# priority-flow-control mode on
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# ets mode on
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# service-policy input type network-qos p5
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# trust-map dot1p t1
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# end
```

LLFC configuration — VLT peer 2

Instead of PFC, you can configure LLFC as follows:

```
OS10# configure terminal
OS10(config)# class-map type network-qos llfc
OS10(config-cmap-nqos)# match qos-group 0-7
OS10(config-cmap-nqos)# exit
OS10(config)# policy-map type network-qos llfc
OS10(config-pmap-network-qos)# class llfc
OS10(config-pmap-c-nqos)# pause buffer-size 50 pause-threshold 30 resume-threshold 10
OS10(config-pmap-c-nqos)# end
OS10# configure terminal
OS10(config)# interface range ethernet 1/1/1,1/1/20,1/1/11,1/1/12
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# flowcontrol transmit on
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# flowcontrol receive on
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# service-policy input type network-qos llfc
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# end
```

WRED/ECN configuration — VLT peer 2

```
OS10# configure terminal
OS10(config)# wred w1
OS10(config-wred)# random-detect ecn
OS10(config-wred)# random-detect color green minimum-threshold 100 maximum-threshold 500 drop-
probability 100
OS10(config-wred)# random-detect color yellow minimum-threshold 100 maximum-threshold 500 drop-
probability 100
OS10(config-wred)# random-detect color red minimum-threshold 100 maximum-threshold 500 drop-
probability 100
OS10(config-wred)# exit
OS10(config)# class-map type queuing cq
OS10(config-cmap-queuing)# match queue 5
OS10(config-cmap-queuing)# exit
OS10(config)# policy-map type queuing pq
OS10(config-pmap-queuing)# class cq
OS10(config-pmap-c-que)# random-detect w1
OS10(config-pmap-c-que)# end
OS10# configure terminal
OS10(config)# interface range ethernet 1/1/1,1/1/20,1/1/11,1/1/12
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# flowcontrol receive off
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# priority-flow-control mode on
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# ets mode on
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# service-policy input type network-qos p5
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# service-policy output type queuing pq
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# trust-map dot1p t1
OS10(conf-range-eth1/1/1,1/1/20,1/1/11,1/1/12)# end
```

Enable DCBx — VLT peer 2

```
OS10# configure terminal
OS10(config)# dcbx enable
```

Configuration on ToR device

System configuration — ToR device

```
NOS# configure terminal
NOS(config)# interface vlan 200
NOS(conf-if-vl-200)# no shutdown
NOS(conf-if-vl-200)# exit
NOS(config)# interface port-channel 2
NOS(conf-if-po-2)# no shutdown
NOS(conf-if-po-2)# exit
NOS(config)# interface range ethernet 1/1/1,1/1/2
NOS(conf-range-eth1/1/1,1/1/2)# channel-group 2 mode active
NOS(conf-range-eth1/1/1,1/1/2)# end
NOS#
NOS# configure terminal
NOS(config)# interface ethernet 1/1/3
NOS(conf-if-eth1/1/3)# switchport mode trunk
NOS(conf-if-eth1/1/3)# switchport trunk allowed vlan 200
NOS(conf-if-eth1/1/3)# end
NOS#
NOS# configure terminal
NOS(config)# interface port-channel 2
NOS(conf-if-po-2)# switchport mode trunk
NOS(conf-if-po-2)# switchport trunk allowed vlan 200
NOS(conf-if-po-2)# end
```

PFC configuration — ToR device

```
NOS# configure terminal
NOS(config)# trust dot1p-map t1
NOS(config-tmap-dot1p-map)# qos-group 0 dot1p 0
NOS(config-tmap-dot1p-map)# qos-group 1 dot1p 1
NOS(config-tmap-dot1p-map)# qos-group 2 dot1p 2
NOS(config-tmap-dot1p-map)# qos-group 3 dot1p 3
NOS(config-tmap-dot1p-map)# qos-group 4 dot1p 4
NOS(config-tmap-dot1p-map)# qos-group 5 dot1p 5
NOS(config-tmap-dot1p-map)# qos-group 6 dot1p 6
NOS(config-tmap-dot1p-map)# qos-group 7 dot1p 7
NOS(config-tmap-dot1p-map)# configure terminal
NOS(config)# class-map type network-qos pfc5
NOS(config-cmap-nqos)# match qos-group 5
NOS(config-cmap-nqos)# exit
NOS(config)# policy-map type network-qos policy5
NOS(config-pmap-network-qos)# class pfc5
NOS(config-pmap-c-nqos)# pause
NOS(config-pmap-c-nqos)# pfc-cos 5
NOS(config-pmap-c-nqos)# end
NOS#
NOS# configure terminal
NOS(config)# interface range ethernet 1/1/1,1/1/2,1/1/3
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# flowcontrol receive off
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# service-policy input type network-qos policy5
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# trust-map dot1p t1
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# priority-flow-control mode on
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# ets mode on
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# end
```

LLFC configuration — ToR device

Instead of PFC, you can configure LLFC as follows:

```
NOS# configure terminal
NOS(config)# class-map type network-qos llfc
NOS(config-cmap-nqos)# match qos-group 0-7
NOS(config-cmap-nqos)# exit
NOS(config)# policy-map type network-qos llfc
NOS(config-pmap-network-qos)# class llfc
NOS(config-pmap-c-nqos)# pause buffer-size 100 pause-threshold 50 resume-threshold 10
NOS(config-pmap-c-nqos)# end
NOS# configure terminal
NOS(config)# interface range ethernet 1/1/1,1/1/2,1/1/3
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# flowcontrol transmit on
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# flowcontrol receive on
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# service-policy input type network-qos llfc
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# end
```

WRED/ECN configuration — ToR device

```
NOS# configure terminal
NOS(config)# wred w1
NOS(config-wred)# random-detect ecn
NOS(config-wred)# random-detect color green minimum-threshold 100 maximum-threshold 500 drop-
probability 100
NOS(config-wred)# random-detect color yellow minimum-threshold 100 maximum-threshold 500 drop-
probability 100
NOS(config-wred)# random-detect color red minimum-threshold 100 maximum-threshold 500 drop-
probability 100
NOS(config-wred)# exit
NOS(config)# class-map type queuing cq
NOS(config-cmap-queuing)# match queue 5
NOS(config-cmap-queuing)# exit
NOS(config)# policy-map type queuing pq
NOS(config-pmap-queuing)# class cq
NOS(config-pmap-c-que)# random-detect w1
NOS(config-pmap-c-que)# end
NOS# configure terminal
NOS(config)# interface range ethernet 1/1/1,1/1/2,1/1/3
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# flowcontrol receive off
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# priority-flow-control mode on
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# ets mode on
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# service-policy input type network-qos policy5
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# service-policy output type queuing pq
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# trust-map dot1p t1
NOS(conf-range-eth1/1/1,1/1/2,1/1/3)# end
```

Enable DCBx — ToR device

```
OS10# configure terminal
OS10(config)# dcbx enable
```

Buffer statistics tracking

OS10 offers the Buffer Statistics Tracking (BST) feature to observe buffer usage across the switch without any impact to performance. This feature maintains separate sets of counters for buffer usage accounting:

- Ingress priority-group
- Ingress service-pool
- Ingress shared-headroom-pool
- Egress queue
- Egress service-pool

You can obtain a snapshot of the buffer statistics for the different buffer objects, such as a snapshot of all ingress priority-groups associated to a port, all egress unicast queues bound to a port, and so on.

You can enable BST at the global level. OS10 tracks buffer utilization and provides the maximum peak statistics value over a period of time and the current value of the monitored BST counter.

Use the `buffer-statistics-tracking` command in SYSTEM-QOS mode to enable BST:

```
OS10# configure terminal
OS10(config)# configure system-qos
OS10(config-sys-qos)# buffer-statistics-tracking
```

Clear the counter

You can choose to reset the peak buffer utilization value and determine a new peak buffer utilization value. Use the `clear qos statistics type buffer-statistics-tracking` command to clear the tracked value and to refresh this counter.

BST tracks peak buffer utilization over a period of time. At any given point in time, the peak buffer usage from the past is displayed.

For example, if you enable BST at time T0 and use the `show` command to view the peak buffer utilization value at time T1, the peak usage between T0 and T1 is displayed. If you view the peak buffer utilization again at time T2, the peak usage between T0 and T2 is displayed. However, if you clear the counter using the `clear qos statistics type buffer-statistics-tracking` command at time T3 and view the peak buffer utilization at time T4, the peak usage between T3 and T4 is displayed.

NOTE: When BST is enabled, if you make any configuration changes that affect the priority group or priority mapping configuration, such as removal of class map, addition of class map to policy map (nqos), and so on, be sure to clear the buffer statistics using the `clear qos statistics type buffer-statistics-tracking` command to view the actual peak buffer utilization for the current configuration.

Advantages of BST include:

- Detecting microburst congestions
- Monitoring buffer utilization and historical trends
- Determining optimal sizes and thresholds for the ingress or egress shared buffers and headroom on a given port or queue based on real-time data

NOTE: BST is not supported on the S4248F-ON and S5148F-ON platforms.

After you disable BST, be sure to clear the counter using the `clear qos statistics type buffer-statistics-tracking` command.

QoS commands

bandwidth

Assigns a percentage of weight to the queue.

Syntax	`bandwidth percent` *value*
Parameters	`percent` *value* — Enter the percentage assignment of bandwidth to the queue, from 1 to 100.
Default	Not configured
Command Mode	POLICY-MAP QUEUE
Usage Information	If you configure this command, you cannot use the `priority` command for the class.

Example	`OS10(conf-pmap-que)# bandwidth percent 70`
Supported Releases	10.2.0E or later

buffer-statistics-tracking

Enables or disables buffer statistics tracking feature globally.

Syntax	`buffer-statistics-tracking`
Parameters	None
Default	Disabled
Command Mode	SYSTEM-QOS
Usage Information	The `no` form of the command disables buffer statistics tracking feature globally. After you disable BST, be sure to clear the counter using the `clear qos statistics type buffer-statistics-tracking` command.

Example	`OS10# configure terminal`
	`OS10(config)# system qos`
	`OS10(config-sys-qos)# buffer-statistics-tracking`

Supported Releases	10.4.3.0 or later

class

Creates a QoS class for a type of policy-map.

Syntax	`class` *`class-name`*
Parameters	*`class-name`* — Enter a name for the class-map. A maximum of 32 characters.
Default	Not configured
Command Mode	POLICY-MAP-QUEUEING POLICY-MAP-QOS POLICY-MAP-NQOS POLICY-MAP-CP POLICY-MAP-APPLICATION
Usage Information	If you define a class-map under a policy-map, the `qos`, `queuing`, or `control-plane` type is the same as the policy-map. You must create this map in advance. The only exception to this rule is when the policy-map type is `trust`, where the class type must be `qos`.
Example	`OS10(conf-pmap-qos)# class c1`
Supported Releases	10.2.0E or later

class-map

Creates a QoS class-map that filters traffic to match packets to the corresponding policy created for your network.

Syntax

```
class-map [type {qos | queuing | control-plane}] [{match-any | match-all}]
class-map-name
```

Parameters

- `type` — Enter a class-map type.
- `qos` — Enter a `qos` type class-map.
- `queuing` — Enter a `queueing` type class-map.
- `control-plane` — Enter a `control-plane` type class-map.
- `match-all` — Determines how packets are evaluated when multiple match criteria exist. Enter the keyword to determine that all packets must meet the match criteria to be assigned to a class.
- `match-any` — Determines how packets are evaluated when multiple match criteria exist. Enter the keyword to determine that packets must meet at least one of the match criteria to be assigned to a class.
- *`class-map-name`* — Enter a class-map name. A maximum of 32 characters.

Defaults

- `qos` — class-map type
- `match-any` — class-map filter

Command Mode

CLASS-MAP-QOS

Usage Information

Apply `match-any` or `match-all` class-map filters to `control-plane`, `qos`, and `queuing` type class-maps.

Example

```
OS10(config)# class-map type qos match-all c1
OS10(conf-cmap-qos)#
```

Command History

10.2.0E or later

clear interface

Clears the statistics per-port or for all ports.

Syntax

```
clear interface [interface node/slot/port[:subport]]
```

Parameters

- *`interface`* — (Optional) Enter the interface type.
- *`node/slot/port[:subport]`* — (Optional) Enter the port information.

Default

Not configured

Command Mode

EXEC

Usage Information

None

Example

```
OS10# clear interface ethernet 1/1/1
```

Supported Releases

10.3.0E or later

clear qos statistics

Clears all QoS-related statistics in the system.

Syntax	`clear qos statistics`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	None
Example	`OS10# clear qos statistics`
Supported Releases	10.2.0E or later

clear qos statistics type

Clears all queue counters for control-plane, qos, and queueing.

Syntax	`clear qos statistics type {{qos | queuing | control-plane | buffer-statistics-tracking} [interface ethernet` *node/slot/port[:subport]*`]}`
Parameters	• `qos`—Clears `qos` type statistics. • `queuing`—Clears `queueing` type statistics. • `control-plane`—Clears `control-plane` type statistics. • `buffer-statistics-tracking`—Clears the peak buffer usage count statistics on all interfaces and service pools. • `interface ethernet` *node-id/slot/port-id [:subport]* — Clears QoS statistics for an Ethernet interface configured for `qos`, `queuing`, or `control-plane`.
Default	Not configured
Command Mode	EXEC
Usage Information	None
Example	`OS10# clear qos statistics type qos interface ethernet 1/1/5`
Example (control-plane)	`OS10# clear qos statistics type control-plane interface ethernet 1/1/7`
Example (queuing)	`OS10# clear qos statistics type queuing interface ethernet 1/1/2`
Example (BST)	`OS10# clear qos statistics type buffer-statistics-tracking`
Supported Releases	10.2.0E or later

control-plane

Enters CONTROL-PLANE mode.

Syntax	`control-plane`
Parameters	None
Default	Not configured
Command Mode	CONTROL-PLANE
Usage Information	If you attach an access-list to the `class-map` type of control-plane, the access-list ignores the `permit` and `deny` keywords.
Example (class-map)	`OS10(config)# class-map type control-plane match-any c1` `OS10(conf-cmap-control-plane)#`
Example (policy-map)	`OS10(config)# policy-map type control-plane p1` `OS10(conf-pmap-control-plane)#`
Supported Releases	10.2.0E or later

control-plane-buffer-size

Configures the buffer size for the CPU pool.

Syntax	`control-plane-buffer-size` *`size-of-buffer-pool`*
Parameters	*`size-of-buffer-pool`*—Enter the buffer size in KB, from 620 KB to 900 KB.
Default	None
Command Mode	SYSTEM-QOS
Usage Information	This command configures the buffer size of the CPU pool. The system allocates a buffer size for the CPU pool from the total system buffer. A minimum guaranteed buffer is allocated for each of the CPU queues and the rest is available for shared usage. The size of the buffer pool varies based on the number of CPU queues and buffer usage by each queue, but it cannot be less than the aggregate of the minimum guaranteed buffer allocated for each of the CPU queues. The `no` version of this command removes the buffer size configured for the CPU pool and returns the buffer size to the default value, 620 KB.
Example	`OS10(config-sys-qos)# control-plane-buffer-size 900`
Supported Releases	10.4.2.0 and later

flowcontrol

Enables or disables link-level flow control on an interface.

Syntax	`flowcontrol [receive | transmit] [on | off]`
Parameters	• `receive` — (Optional) Indicates the port can receive flow control packets from a remote device.

> **NOTE: In S5148F-ON, when receive is turned on, it enables decoding of both LLFC and PFC frames on that port.**

- `transmit` — (Optional) Indicates the local port can send flow control packets to a remote device.
- `on` — (Optional) When used with `receive`, allows the local port to receive flow control traffic. When used with `transmit`, allows the local port to send flow control traffic to the remote device.
- `off` — (Optional) When used with `receive`, disables the remote device from sending flow control traffic to the local port. When used with `transmit`, disables the local port from sending flow control traffic to the remote device.

Default	Disabled (`off`)
Command Mode	INTERFACE
Usage Information	The `no` version of this command returns the value to the default.
Example	`OS10(conf-if-eth1/1/2)# flowcontrol transmit on`
Supported Releases	10.3.0E or later

hardware deep-buffer-mode

Configures Deep Buffer mode.

Syntax	`hardware deep-buffer-mode`
Parameters	None
Defaults	Disabled
Command Modes	CONFIGURATION
Usage Information	Deep Buffer mode configuration takes effect only after you save it in the startup configuration and reboot the switch. The `no` version of this command disables Deep Buffer mode.
Example	`OS10(config)# hardware deep-buffer-mode`
Supported Releases	10.4.3.0 or later

match

Configures match criteria for the QoS policy.

Syntax

```
match {cos cos-number | ip [access-group name name | dscp dscp-value |
precedence value] | ipv6 [access-group name name [set dscp dscp-value]] | mac
access-group acl-name | not [ip | cos] vlan vlan-id} [set dscp dscp-value]
```

Parameters

- `cos cos–number` — Enter a queue number for the CoS match criteria, from 0 to 7.
- `ip` — Enter the IPv4 match criteria.
- `access-group name name` — (Optional) Enter the IPv4 access-group name.
- `dscp dscp-value` — (Optional) Enter a DSCP value for L3 DSCP match criteria, from 0 to 63.
- `precedence value` — (Optional) Enter a precedence value for L3 precedence match criteria, from 0 to 7.
- `ipv6` — Enter the IPv6 match criteria.

- `access-group name` *name* — (Optional) Enter the IPv6 access-group name.
- `set dscp` *dscp-value* — (Optional) Configure a DSCP value for L3 DSCP match criteria, from 0 to 63.
- `mac access-group name` *name* — Enter an access-group name for the MAC access-list match criteria. A maximum of 140 characters.
- `set dscp` *dscp-value* — Enter a DSCP value for marking the DSCP packets, from 0 to 63.
- `not` — Enter the IP or CoS to negate the match criteria.
- `vlan` *vlan-id* — Enter a VLAN number for VLAN match criteria, from 1 to 4093.

Default Not configured

Command Mode CLASS-MAP

Usage Information In a `match-any` class, you can enter multiple match criteria. In a `match-all` class, if the match case is `access-group`, no other match criteria is allowed. If you attach the access-list to `class-map type control–plane`, the access-list ignores the `permit` and `deny` keywords.

Example

```
OS10(conf-cmap-qos)# match ip access-group name ag1
OS10(config-cmap-qos)# match ipv6 access-group name ACLv6 set dscp 40
```

Supported Releases 10.2.0E or later

match cos

Matches a cost of service (CoS) value to L2 dot1p packets.

Syntax `match [not] cos` *cos-value*

Parameters

- *cos-value* — Enter a CoS value, from 0 to 7.
- `not` — Enter `not` to cancel the match criteria.

Default Not configured

Command Modes CLASS-MAP

Usage Information You cannot have two match statements with the same filter-type. If you enter two match statements with the same filter-type, the second statement overwrites the first statement.

Example

```
OS10(conf-cmap-qos)# match cos 3
```

Supported Releases 10.2.0E or later

match dscp

Configures a DSCP value as a match criteria for a class-map.

Syntax `match [not] {ip | ipv6 | ip-any } dscp [`*dscp-list* `|` *dscp-list*`]`

Parameters

- `not` — (Optional) Enter to cancel a previously applied match criteria.
- `ip` — Enter to use IPv4 as the match protocol.
- `ipv6` — Enter to use IPv6 as the match protocol.
- `ip-any` — Enter to use both IPv4 and IPv6 as the match protocol.

- `dscp` *`dscp-list | dscp-list`* — Enter a DSCP value in single numbers, comma separated, or a hyphenated range, from 0 to 63.

Default Not configured

Command Mode CLASS-MAP

Usage Information You cannot enter two match statements with the same filter-type. If you enter two match statements with the same filter-type, the second statement overwrites the first statement. The `match-all` option in a class-map does not support `ip-any`. Select either `ip` or `IPv6` for the `match-all` criteria. If you select `ip-any`, you cannot select `ip` or `ipv6` for the same filter type.

Example

```
OS10(conf-cmap-qos)# match ip-any dscp 17-20
```

Supported Releases 10.2.0E or later

match precedence

Configures IP precedence values as a match criteria.

Syntax

```
match [not] {ip | ipv6 | ip-any} precedence precedence-list
```

Parameters

- `not` — Enter to cancel a previously applied match precedence rule.
- `ip` — Enter to use IPv4 as the match precedence rule.
- `ipv6` — Enter to use IPv6 as the match precedence rule.
- `ip-any` — Enter to use both IPv4 and IPv6 as the match precedence rule.
- `precedence` *`precendence-list`* — Enter a precedence-list value, from 0 to 7.

Default Not configured

Command Mode CLASS-MAP

Usage Information You cannot enter two match statements with the same filter-type. If you enter two match statements with the same filter-type, the second statement overwrites the first statement.

Example

```
OS10(conf-cmap-qos)# match not ipv6 precedence 3
```

Supported Releases 10.2.0E or later

match queue

Configures a match criteria for a queue.

Syntax

```
match queue queue-number
```

Parameters *`queue-number`* — Enter a queue number, from 0 to 7.

Default Not configured

Command Mode CLASS-MAP

Usage Information You can configure this command only when the class-map type is `queuing`. You cannot enter two match statements with the same filter-type. If you enter two match statements with the same filter-type, the second statement overwrites the first statement.

Example	`OS10(conf-cmap-queuing)# match queue 1`
Supported Releases	10.2.0E or later

match vlan

Configures a match criteria based on the VLAN ID number.

Syntax	`match vlan vlan-id`
Parameters	*vlan-id* — Enter a VLAN ID number, from 1 to 4093.
Default	Not configured
Command Mode	CLASS-MAP
Usage Information	You cannot enter two match statements with the same filter-type. If you enter two match statements with the same filter-type, the second statement overwrites the first statement.
Example	`OS10(conf-cmap-qos)# match vlan 100`
Supported Releases	10.2.0E or later

mtu

Calculates the buffer size allocation for matched flows.

Syntax	`mtu size`
Parameters	*size* — Enter the size of the buffer (1500 to 9216).
Default	9216
Command Mode	POLICY-MAP-CLASS-MAP
Usage Information	The `no` version of this command returns the value to the default.
Example	`OS10(conf-pmap-nqos-c)# mtu 2500`
Supported Releases	10.3.0E or later

pause

Enables a pause based on buffer limits for the port to start or stop communication to the peer.

Syntax	`pause [buffer-size size pause-threshold xoff-size resume-threshold xon-size]`
Parameters	• `buffer-size size` — (Optional) Enter the ingress buffer size used as a guaranteed buffer in KB, . – Default values for PFC: 10G, 25G–183KB, 40G–375KB, 100G–446KB – Default values for LLFC: 10G,25G–207.5KB, 40G,100G–300.5KB • `pause-threshold xoff-size` — (Optional) Enter the buffer limit for the port to start or initiate a `pause` to the peer in KB, . – Default values for PFC: 10G, 25G–96KB, 40G–192KB, 100G–232KB – Default values for LLFC: 10G,25G–198.5KB, 40G,100G–264.5KB

- `resume-threshold` *`xon-size`* — (Optional) Enter the buffer limit for the port to stop or cancel sending a `pause` to the peer in KB .
  - Default values for PFC: 10G, 25G–87KB, 40G–183KB, 100G–214KB
  - Default values for LLFC: 10G,25G–9KB, 40G,100G–36KB

Default See parameter values

Command Mode POLICY-MAP-CLASS-MAP

Usage Information Only use this command under the `network-qos` policy type. Buffer-size, pause-thresholds, and resume-thresholds vary based on platform. Add the policy-map with `pause` to system-qos to service an input to enable `pause` on all ports, based on a per-port link-level Flow-Control mode. The `xoff` and `xon` threshold settings for link-level flow-control are applied on ports where all traffic classes must be mapped to a single PG. Platform-specific default values are based on MTU sizes of 9216 and cable length of 100 meters. The `no` version of this command returns the value to the default.

Example

```
OS10(conf-pmap-c-nqos)# pause buffer-size 45 pause-threshold 25 resume-
threshold 10
```

Example (global and shared buffer)

```
OS10(config)# policy-map type network-qos nqGlobalpolicy1
OS10(conf-cmap-nqos)# class CLASS-NAME
OS10(conf-cmap-nqos-c)# pause buffer-size 45 pause-threshold 30 resume-
threshold 30
```

```
OS10(config)# policy-map type network-qos nqGlobalpolicy1
OS10(conf-cmap-nqos)# class type network-qos nqclass1
OS10(conf-cmap-nqos-c)# pause buffer-size 45 pause-threshold 30 resume-
threshold 10
```

Supported Releases 10.3.0E or later

pfc-cos

Configures priority flow-control for cost of service (CoS).

Syntax `pfc-cos` *`cos-value`*

Parameters *`cos-value`* — Enter a single, comma-delimited, or hyphenated range of CoS values for priority flow-control to enable, from 0 to 7.

NOTE: The range 0-7 is invalid. All other ranges, including 0-6 and 1-7 are valid.

Default Not configured

Command Mode POLICY-MAP-CLASS-MAP

Usage Information To configure link-level flow-control, do not configure `pfc-cos` for the matched class for this policy. Add the policy-map with the `pfc-cos` configuration to `system-qos` to service an input to enable priority flow-control behavior on all ports, based on a per-port Priority Flow-Control Enable mode. Add the policy-map with the `pfc-cos` configuration to interface configurations to service at input and enable Priority Flow-Control on that particular port, based on the port's Priority Flow-Control Enable mode. If you configure 40G to 10G mode on interfaces and `pause` (no drop) is enabled on `system-qos`, all queues may or may not drop traffic based on the availability of buffers. The `no` version of this command returns the value to the default.

Example

```
OS10(conf-pmap-c-nqos)# pfc-cos 0-2
```

Example (global buffer/shared buffer)	`OS10(config)# policy-map type network-qos nqGlobalpolicy1`
	`OS10(conf-cmap-nqos)# class CLASS-NAME`
	`OS10(conf-cmap-nqos-c)# pause buffer-size 45 pause-threshold 25 resume-threshold 10`
	`OS10(conf-cmap-nqos-c)# pfc-cos 0-2`
	`OS10(conf-cmap-nqos-c)# queue-limit 140`

Supported Releases	10.3.0E or later

pfc-max-buffer-size

Configures the maximum buffer size for priority flow-control enabled flows.

Syntax	`pfc-max-buffer-size` *`max-buffer-size`*
Parameters	*`max-buffer-size`* — Enter the maximum buffer size in KB.
Default	None
Command Mode	SYSTEM-QOS
Usage Information	This command configures the maximum size of the lossless buffer pool. The `no` version of this command removes the maximum buffer size limit.
Example	`OS10(config-sys-qos)# pfc-max-buffer-size 2000`
Supported Releases	10.4.0E(R1) or later

pfc-shared-buffer-size

Changes the shared buffers size limit for priority flow-control enabled flows.

Syntax	`pfc-shared-buffer-size` *`buffer-size`*
Parameters	*`buffer-size`* — Enter the size of the priority flow-control buffer in KB, from 0 to 8911.
Default	832 KB
Command Mode	SYSTEM-QOS
Usage Information	The `no` version of this command returns the value to the default.
Example	`OS10(conf-sys-qos)# pfc-shared-buffer-size 2000`
Supported Releases	10.3.0E or later

pfc-shared-headroom-buffer-size

Configures the shared headroom size for absorbing the packets after pause frames generate.

NOTE: This command is available only on the following platforms:

- S5212F-ON, S5224F-ON, S5232F-ON, S5248F-ON, S5296F-ON
- Z9100-ON
- Z9264F-ON

Syntax	`pfc-shared-headroom-buffer-size` *`headroom-buffer-size`*
Parameters	*`headroom-buffer-size`* — Enter the size of the priority flow-control headroom buffer in KB, from 1 to 3399.
Default	1024 KB
Command Mode	SYSTEM-QOS
Usage Information	All PFC-enabled priority groups can use the shared headroom space. Headroom is the buffer space that absorbs the incoming packets after the PFC frames reach the sender. After the threshold is reached, PFC frames generate towards the sender. The packets sent by the sender after the PFC frames generate are absorbed into the Headroom buffer. The `no` version of this command returns the value to the default.
Example	`OS10(conf-sys-qos)# pfc-shared-headroom-buffer-size 2000`
Supported Releases	10.4.0E(R1) or later

police

Configures traffic policing on incoming traffic.

Syntax	`police {cir` *`committed-rate`* `[bc` *`committed-burst-size`*`]} {pir` *`peak-rate`* `[be` *`peak-burst-size`*`]}`
Parameters	• `cir` *`committed-rate`* — Enter a committed rate value in kilo bits per second, from 0 to 4000000. • `bc` *`committed-burst-size`* — (Optional) Enter the committed burst size in packets for control plane policing and in KB for data packets, from 16 to 200000. • `pir` *`peak-rate`* — Enter a peak-rate value in kilo bits per second, from 0 to 40000000. • `be` *`peak-burst-size`* — (Optional) Enter a peak burst size in kilo bytes, from 16 to 200000.
Defaults	• `bc` *`committed-burst-size`* value is 200 KB for `control plane` and 100 KB for all other class-map types • `be` *`peak-burst-size`* value is 200 KB for `control plane` and 100 KB for all other class-map types
Command Mode	POLICY-MAP-CLASS-MAP
Usage Information	If you do not provide the peak-rate `pir` values, the committed-rate `cir` values are taken as the `pir` values. Only the ingress QoS policy type supports this command. For control-plane policing, the rate values are in pps.
Example	`OS10(conf-pmap-c-qos)# police cir 5 bc 30 pir 20 be 40`
Supported Releases	10.2.0E or later

policy-map

Enters QoS POLICY-MAP mode and creates or modifies a QoS policy-map.

Syntax	`policy-map` *`policy-map-name`* `[type {qos | queuing | control-plane | application | network-qos }]`
Parameters	• *`policy-map-name`* — Enter a class name for the policy-map. A maximum of 32 characters. • `type` — Enter the policy-map type. – `qos` — Create a `qos` policy-map type.

- queuing — Create a queueing policy-map type.
- control-plane — Create a control-plane policy-map type.
- application — Create an application policy-map type.
- network-qos — Create a network-qos policy-map type.

Defaults qos = class-map type and match-any = class-map filter

Command Mode CONFIGURATION

Usage Information The no version of this command deletes a policy-map.

Example
```
OS10(config)# policy-map p1
```

Example (Queuing)
```
OS10(config)# policy-map type queuing p1
```

Supported Releases 10.2.0E or later

priority

Sets the scheduler as a strict priority.

Syntax priority

Parameters None

Default WDRR — when priority is mentioned, it moves to SP with default level 1.

Command Mode POLICY-MAP-CLASS-MAP

Usage Information If you use this command, bandwidth is not allowed. Only the egress QoS policy type supports this command.

Example
```
OS10(conf-pmap-que)# priority
```

Supported Releases 10.2.0E or later

priority-flow-control mode

Enables or disables Priority Flow-Control mode on an interface.

Syntax priority-flow-control mode [on]

Parameters

- on — (Optional) Enables Priority Flow-Control mode.

Default Disabled

Command Mode INTERFACE

Usage Information Before enabling priority flow-control on a interface, verify a matching network-qos type policy is configured with the pfc-cos value for an interface. Use this command to disable priority flow-control if you are not using a network-qos type policy for an interface. The no version of this command returns the value to the default.

Example
```
OS10(conf-if-eth1/1/2)# priority-flow-control mode on
```

Supported Releases 10.3.0E or later

qos-group dot1p

Configures a dot1p trust map to the traffic class.

Syntax `qos-group tc-list [dot1p values]`

Parameters

- `qos-group tc-list` — Enter the traffic single value class ID, from 0 to 7.
- `dot1p values` — (Optional) Enter either single, comma-delimited, or a hyphenated range of dot1p values, from 0 to 7.

Default 0

Command Mode TRUST-MAP

Usage Information If the trust map does not define dot1p values to any traffic class, those flows map to the default traffic class `0`. If some of the dot1p values are already mapped to an existing traffic class, you see an error. You must have a 1:1 dot1p-to-traffic class mapping for PFC-enabled CoS values. You must also have a common dot1p trust map for all interfaces using DCB. The `no` version of this command returns the value to the default.

Example

```
OS10(conf-tmap-dot1p-qos)# qos-group 5 dot1p 5
```

Supported Releases 10.3.0E or later

qos-group dscp

Configures a DSCP trust map to the traffic class.

Syntax `qos-group tc-list [dscp values]`

Parameters

- `qos-group tc-list` — Enter the traffic single value class ID, from 0 to 7.
- `dscp values` — (Optional) Enter either single, comma-delimited, or a hyphenated range of DSCP values, from 0 to 63.

Default 0

Command Mode TRUST-MAP

Usage Information If the trust map does not define DSCP values to any traffic class, those flows map to the default traffic class `0`. If some of the DSCP values are already mapped to an existing traffic class, you will see an error. The `no` version of this command returns the value to the default.

Example

```
OS10(conf-tmap-dscp-qos)# qos-group 5 dscp 42
```

Supported Releases 10.3.0E or later

qos-rate-adjust

Configures additional number of data bytes to add to overhead fields per frame for rate calculations.

Syntax `qos-rate-adjust [value-of-adjust]`

Parameters *value-of-adjust*—Number of bytes to add to overhead fields in each frame, from 1 to 31.

Default 0

Command Mode CONFIGURATION

Usage Information The `no` form of this command removes the rate adjustment configuration and is the same as using the `qos-rate-adjust 0` command.

Example

```
OS10(config)# qos-rate-adjust 10
```

Supported Releases 10.4.3.0 or later

queue-limit

Configures static or dynamic shared buffer thresholds.

Syntax `queue-limit {queue-len value | thresh-mode [dynamic threshold-alpha-value | static threshold-value]}`

Parameters

- `queue-len value` — Enter the guaranteed size for the queue, from 0 to 8911.
  - 45 KB (10G)/111 KB (40G) if the queue is priority flow control enabled
  - 2 KB (10G)/8 KB (40G) if the queue is lossy/link-level flow control
  - If this is a priority flow-control queue, this configuration is invalid
  - Only supported for POLICY-MAP-CLASS-MAP (`pmap-c-queue`) mode
- `thresh-mode` — (Optional) Buffer threshold mode.
- `dynamic thresh-alpha-value` — (Optional) Enter the value indexes to calculate the shared threshold to the enabled dynamic shared buffer threshold, from 0 to 10. Defaults:
  - 0 = 1/128
  - 1 = 1/64
  - 2 = 1/32
  - 3 = 1/16
  - 4 = 1/8
  - 5 = 1/4
  - 6 = 1/2
  - 7 = 1
  - 8 = 2
  - 9 = 4
  - 10 = 8
- `static thresh-value` — (Optional) Enter the static shared buffer threshold value in Bytes, from 1 to 65535.

Default Not configured

Command Mode POLICY-MAP-CLASS-MAP

Usage Information Use the `queue-len value` parameter to set the minimum guaranteed queue length for a queue. The `no` version of this command returns the value to the default.

Example

```
OS10(config)# policy-map type network-qos nqGlobalpolicy1
OS10(conf-cmap-nqos)# class type network-qos nqclass1
OS10(conf-cmap-nqos-c)# pause buffer-size 45 pause-threshold 30 resume-
threshold 10
OS10(conf-cmap-nqos-c)# queue-limit 150
```

Example (queue)

```
OS10(config)# policy-map type queuing pmap1
OS10(config-pmap-queuing)# class cmap1
OS10(config-pmap-c-que)# queue-limit queue-len 100
OS10(config-pmap-c-que)# queue-limit thresh-mode static 50
```

Supported Releases 10.3.0E or later

queue bandwidth

Configures a bandwidth for a given queue on interface.

Syntax `queue queue-number bandwidth bandwidth-percentage`

Parameters

- *queue-number* — Enter the queue number.
- *bandwidth-percentage* — Enter the percentage of bandwidth.

Default Not configured

Command Mode POLICY-MAP-CLASS-MAP

Usage Information The `no` version of this command removes the bandwidth from the queue.

Example None

Supported Releases 10.4.0E(R1) or later

queue qos-group

Configures a dot1p traffic class to a queue.

Syntax `queue number [qos-group dot1p-values]`

Parameters

- `queue` *number* — Enter the traffic single value queue ID, from 0 to 7.
- `qos-group` *dot1p-values* — (Optional) Enter either single, comma-delimited, or a hyphenated range of dot1p values, from 0 to 7.

Default 0

Command Mode TRUST-MAP

Usage Information If the trust map does not define traffic class values to a queue, those flows map to the default queue `0`. If some of the traffic class values are already mapped to an existing queue, you see an error. The `no` version of this command returns the value to the default.

Example

```
OS10(conf-tmap-tc-queue-qos)# queue 2 qos-group 5
```

Supported Releases 10.3.0E or later

random-detect (interface)

Assigns a WRED profile to the specified interface.

Syntax `random-detect wred-profile`

Parameters *wred-profile* — Enter the name of an existing WRED profile.

Default Not configured

Command Mode INTERFACE

Usage Information The `no` version of this command removes the WRED profile from the interface.

Example

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# random-detect test_wred
```

Supported Releases 10.4.0E(R1) or later

random-detect (queue)

Assigns a WRED profile to the specified queue.

Syntax `random-detect` *wred-profile-name*

Parameters *wred-profile-name* — Enter the name of an existing WRED profile.

Default Not configured

Command Mode PMAP-C-QUE

Usage Information The `no` version of this command removes the WRED profile from the queue.

Example

```
OS10(config)# policy-map type queuing p1
OS10(config-pmap-queuing)# class c1
OS10(config-pmap-c-que)# random-detect test_wred
```

Supported Releases 10.4.0E(R1) or later

random-detect color

Configures the threshold of WRED profile for available colors.

Syntax `random-detect color` *color-name* `minimum-threshold` *minimum-value* `maximum-threshold` *maximum-value* `drop-probability` *drop-rate*

Parameters

- *color-name* — Enter the color of drop precedence for the WRED profile. The available options are green, yellow, and red.
- *minimum-value* — Enter the minimum threshold value for the specified color, from 1 to 12480.
- *maximum-value* — Enter the maximum threshold value for the specified color, from 1 to 12480.
- *drop-rate* — Enter the rate of drop precedence in percentage, from 0 to 100.

Default Not configured

Command Mode WRED CONFIGURATION

Usage Information The `no` version of this command removes the WRED profile.

Example

```
OS10(config)# wred test_wred
OS10(config-wred)# random-detect color green minimum-threshold 100 maximum-
threshold 300 drop-probability 40
```

Supported Releases 10.4.0E(R1) or later

random-detect ecn

Enables explicit congestion notification (ECN) for the WRED profile.

Syntax `random-detect ecn`

Parameters None

Default Not configured

Command Mode WRED CONFIGURATION

Usage Information The `no` version of this command disables ECN.

Example

```
OS10(config)# wred test_wred
OS10(config-wred)# random-detect ecn
```

Supported Releases 10.4.0E(R1) or later

random-detect ecn

Enables ECN for the system globally.

Syntax `random-detect ecn`

Default Not configured

Command Mode SYSTEM QOS

Usage Information The `no` version of this command disables ECN globally.

NOTE: This command enables ECN globally and is supported only on the S4200–ON Series platform. In the SYSTEM QOS mode, this command is not available on other platforms. Also, you can configure ECN only per queue; you cannot configure ECN on an interface or service pool on the S4200–ON Series platform.

Example

```
applicableOS10(config)# system-qos
OS10(config-sys-qos)# random-detect ecn
```

Supported Releases 10.4.1.0 or later

random-detect pool

Assigns a WRED profile to the specified global buffer pool.

Syntax `random-detect pool` *`pool-value wred-profile-name`*

Parameters

- *`pool-value`* — Enter the pool value, from 0 to 1.
- *`wred-profile-name`* — Enter the name of an existing WRED profile.

Default Not configured

Command Mode SYSTEM-QOS

Usage Information The `no` version of this command removes the WRED profile from the interface.

Example

```
OS10(config)# system qos
OS10(config-sys-qos)# random-detect pool 0 test_wred
```

Supported Releases 10.4.0E(R1) or later

random-detect weight

Configures the exponential weight value used to calculate the average queue depth for the WRED profile.

Syntax `random-detect weight` *`weight-value`*

Parameters *`weight-value`* — Enter a value for the weight, from 1 to 15.

Default Not configured

Command Mode WRED CONFIGURATION

Usage Information The `no` version of this command removes the weight factor from the WRED profile.

Example

```
OS10(config)# wred test_wred
OS10(config-wred)# random-detect weight 10
```

Supported Releases 10.4.0E(R1) or later

service-policy

Configures the input and output service policies.

Syntax `service-policy {input | output} [type {qos | queuing | network-qos}]` *`policy-map-name`*

Parameters

- `input` — Enter to assign a QoS policy to the interface input.
- `output` — Enter to assign a QoS policy to the interface output.
- `qos` — Enter to assign a `qos` type policy-map.
- `queuing` — Enter to assign the `queuing` type policy-map.
- `network-qos` — Enter to assign the `network-qos` type policy-map.
- *`policy-map-name`* — Enter the policy-map name. A maximum of 32 characters.

Default Not configured

Command Mode INTERFACE

Usage Information Attach only one policy-map to the interface input and output for each `qos` and `queuing` policy-map type. You can attach four service-policies to the system QoS — one each for `qos`, `queueing`, and `network-qos` type policy-maps. When you configure interface-level policies and system-level policies, the interface-level policy takes precedence over the system-level policy.

Example

```
OS10(conf-if-eth1/1/7)# service-policy input type qos p1
```

Supported Releases 10.2.0E or later

set cos

Sets a cost of service (CoS) value to mark L2 802.1p (dot1p) packets.

Syntax	`set cos cos-value`
Parameters	*cos-value* — Enter a CoS value, from 0 to 7.
Default	Not configured
Command Mode	POLICY-MAP-CLASS-MAP
Usage Information	You cannot enter two set statements with the same action-type. If you enter two set statements with the same action-type, the second statement overwrites the first. When class-map type is `qos`, the qos-group corresponds to data queues 0 to 7.
Example	`OS10(conf-pmap-c-qos)# set cos 6`
Supported Releases	10.2.0E or later

set dscp

Sets the drop precedence for incoming packets based on their DSCP value and color map profile.

Syntax	`set dscp dscp-value`
Parameters	*dscp-value* — Enter a DSCP value, from 0 to 63.
Default	Not configured
Command Mode	POLICY-MAP-CLASS-MAP
Usage Information	When class-map type is `qos`, the qos-group corresponds to data queues 0 to 7.
Example	`OS10(conf-pmap-c-qos)# set dscp 10`
Supported Releases	10.2.0E or later

set qos-group

Configures marking for the QoS-group queues.

Syntax	`set qos-group queue-number`
Parameters	*queue-number* — Enter a queue number, from 0 to 7.
Default	Not configured
Command Mode	POLICY-MAP-CLASS-MAP
Usage Information	This command supports only the `qos` or `control-plane` ingress policy type. When the class-map type is `control-plane`, the qos-group corresponds to CPU queues 0 to 11. When the class-map type is `qos`, the qos-group corresponds to data queues 0 to 7.
Example	`OS10(conf-pmap-c-qos)# set qos-group 7`

Supported Releases 10.2.0E or later

shape

Shapes the outgoing traffic rate.

Syntax `shape {min {kbps | mbps} min-value [burst-size]} {max {kbps | mbps} max-value [max-burst-size]}`

Parameters

- `min` — Enter the minimum committed rate in unit in kbps, mbps.
- `kbps` — Enter the committed rate unit in kilobits per second, from 0 to 40000000.
- `mbps` — Enter the committed rate unit in megabits per second, from 0 to 40000.
- *`burst-size`* — Enter the burst size in kilobytes per packet, from 0 to 10000 or 1 to 1073000.
- `max` — Enter the maximum peak rate in kbps, mbps.
- *`max-burst-size`* — Enter the burst size in kilobytes per packets, from 0 to 10000 or 1 to 1073000.

Default Maximum burst size is 50 kb

Command Mode POLICY-MAP-CLASS-MAP

Usage Information This command only supports the ingress QoS policy type. You must enter both the minimum and maximum values. If you enter the rate value in pps, the burst provided is in packets. If you enter the rate in kbps or mbps, the burst is provided in kb.

Example

```
OS10(conf-pmap-c-que)# shape min kbps 11 max kbps 44
```

Supported Releases 10.2.0E or later

show class-map

Displays configuration details of all existing class-maps.

Syntax `show class-map [type {control-plane | qos | queuing | network-qos} class-map-name]`

Parameters

- `type` — Enter the policy-map type — `qos`, `queuing`, or `control-plane`.
- `qos` — Displays all policy-maps of `qos` type.
- `queuing` — Displays all policy-maps of `queuing` type.
- `network-qos` — Displays all policy-maps of `network-qos` type.
- `control-plane` — Displays all policy-maps of `control-plane` type.
- *`class-map-name`* — Displays the QoS class-map name.

Default Not configured

Command Mode EXEC

Usage Information This command displays all class-maps of `qos`, `queuing`, `network-qos`, or `control-plane` type. The *`class-map-name`* parameter displays all details of a configured class-map name.

Example

```
OS10# show class-map type qos c1
    Class-map (qos):  c1 (match-all)
      Match(not): ip-any dscp 10
```

Supported Releases 10.2.0E or later

show control-plane buffers

Displays the pool type, reserved buffer size, and the maximum threshold value for each of the CPU queues.

Syntax `show control-plane buffers`

Parameters None

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show control-plane buffers
queue-number pool-type   rsvd-buf-size   threshold-mode     threshold-value
-------------------------------------------------------------------------------
0            lossy            1664            static
20800
1            lossy            1664            static
20800
2            lossy            1664            static
48880
3            lossy            9216            static
48880
4            lossy            1664            static
20800
5            lossy            1664            static
48880
6            lossy            1664            static
48880
7            lossy            1664            static
48880
8            lossy            1664            static
48880
9            lossy            9216            static
48880
10           lossy            1664            static
48880
11           lossy            1664            static
48880
12           lossy            1664            static
48880
13           lossy            9216            static
48880
14           lossy            1664            static
48880
15           lossy            9216            static
48880
16           lossy            1664            static
48880
17           lossy            1664            static
48880
18           lossy            1664            static
48880
19           lossy            1664            static
48880
20           lossy            1664            static
20800
```

```
21            lossy            1664             static
20800
22            lossy            1664             static           20800
```

Supported Releases 10.4.2 and later

show control-plane buffer-stats

Displays the control plane buffer statistics for each of the CPU queues.

Syntax `show control-plane buffer-stats`

Parameters None

Default A predefined default profile exists.

Command Mode EXEC

Usage Information None

Example

```
OS10# show control-plane buffer-stats
Queue    TX              TX                  Used reserved    Used shared
         pckts           bytes               buffers           buffers
-------------------------------------------------------------------------
0          0               0                   0                 0
1          0               0                   0                 0
2          0               0                   0                 0
3          0               0                   0                 0
4          0               0                   0                 0
5          0               0                   0                 0
6          3               204                 0                 0
7          6               408                 0                 0
8          0               0                   0                 0
9          0               0                   0                 0
10         0               0                   0                 0
11         0               0                   0                 0
12         0               0                   0                 0
13         0               0                   0                 0
14         0               0                   0                 0
15         0               0                   0                 0
16         0               0                   0                 0
17         0               0                   0                 0
18         0               0                   0                 0
19         0               0                   0                 0
20         0               0                   0                 0
21         0               0                   0                 0
22         0               0                   0                 0
```

Supported Releases 10.4.2 and later

show control-plane info

Displays control-plane queue mapping and rate limits.

Syntax `show control-plane info`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information Monitors statistics for the control-plane and to troubleshoot CoPP.

Example

```
OS10# show control-plane info
Queue  Min Rate Limit(in pps) Max Rate Limit(in pps)   Protocols
0      600                    600                    ISCSI UNKNOWN UNICAST
1      1000                   1000                   SFLOW
2      400                    400                    IGMP MLD PIM
3      600                    1000                   VLT NDS
4      500                    1000                   IPV6_ICMP IPV4_ICMP
5      500                    1000                   ICMPV6_RS ICMPV6_NS
ICMPV6_RA ICMPV6_NA
6      500                    1000                   ARP_REQ SERVICEABILITY
7      500                    1000                   ARP_RESP
8      500                    500                    SSH TELNET TACACS NTP FTP
9      600                    600                    FCOE
10     600                    1000                   LACP
11     400                    400                    RSTP PVST MSTP
12     500                    500                    DOT1X LLDP
13     600                    1000                   IPV6_OSPF IPV4_OSPF
14     600                    1000                   OSPF_HELLO
15     600                    1000                   BGP
16     500                    500                    IPV6_DHCP IPV4_DHCP
17     600                    1000                   VRRP
18     700                    700                    BFD
19     700                    1000                   OPEN_FLOW REMOTE CPS
20     300                    300                    MCAST DATA
21     100                    100                    ACL LOGGING
22     300                    300                    MCAST KNOWN DATA
```

Supported Releases 10.2.0E or later

show control-plane statistics

Displays counters of all the CPU queue statistics.

Syntax `show control-plane info`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show control-plane statistics
Queue   Packets   Bytes    Dropped Packets   Dropped Bytes
0       0         0        0                 0
1       0         0        0                 0
2       0         0        0                 0
3       0         0        0                 0
4       0         0        0                 0
5       0         0        0                 0
6       3         204      0                 0
7       6         408      0                 0
8       0         0        0                 0
9       0         0        0                 0
10      0         0        0                 0
11      0         0        0                 0
12      0         0        0                 0
13      0         0        0                 0
14      0         0        0                 0
15      0         0        0                 0
16      0         0        0                 0
17      0         0        0                 0
```

```
18      0          0         0                   0
19      0          0         0                   0
20      0          0         0                   0
21      0          0         0                   0
22      0          0         0                   0
OS10#
```

Supported Releases 10.2.0E or later

show hardware deep-buffer-mode

Displays the status of DeepB buffer mode in the current and next boot of the switch.

Syntax `show hardware deep-buffer-mode`

Parameters None

Defaults Not configured

Command Modes EXEC

Example Example: default setting

```
OS10# show hardware deep-buffer-mode
  Deep Buffer Mode Configuration Status
-----------------------------------------
    Current-boot Settings : Disabled

    Next-boot Settings    : Disabled
```

Example: saved to startup configuration

```
OS10# show hardware deep-buffer-mode
  Deep Buffer Mode Configuration Status
-----------------------------------------
    Current-boot Settings : Disabled

    Next-boot Settings    : Enabled
```

Example: switch reloaded

```
OS10# show hardware deep-buffer-mode
  Deep Buffer Mode Configuration Status
-----------------------------------------
    Current-boot Settings : Enabled

    Next-boot Settings    : Enabled
```

Supported Releases 10.4.3.0 or later

show interface priority-flow-control

Displays the priority flow-control, operational status, CoS bitmap, and statistics per port.

Syntax `show interface ethernet 1/1/1 priority-flow-control [details]`

Parameters `details` — (Optional) Displays all priority flow control information for an interface.

Default Not configured

Command Mode EXEC

Usage Information None

Example (Details)

```
OS10# show interface priority-flow-control details

TenGig 1/1:
Admin Mode: On
OperStatus: On
PFC Priorites: 0,4,7
Total Rx PFC Frames: 300
Total Tx PFC Frames: 200
Cos     Rx      Tx
----------------------
0       0       0
1       0       0
2       0       0
3       300     200
4       0       0
5       0       0
6       0       0
7       0       0
```

Supported Releases 10.3.0E or later

show qos interface

Displays the QoS configuration applied to a specific interface.

Syntax `show qos interface ethernet` *node/slot/port[:subport]*

Parameters *node/slot/port[:subport]* — Enter the Ethernet interface information.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show qos interface ethernet 1/1/10
Ethernet 1/1/10
    unknown-unicast-storm-control : 100 pps
    multicast-storm-control : 200 pps
    broadcast-storm-control : Disabled
    flow-control-rx: Enabled
    flow-control-tx: Disabled
      Service-policy (Input)(qos): p1
```

Supported Releases 10.2.0E or later

show policy-map

Displays information on all existing policy-maps.

Syntax `show policy-map type {control-plane | qos | queuing | network-qos}] [`*policy-map-name*`]`

Parameters

- `type` — Enter the policy-map type — `qos`, `queuing`, or `control-plane`.
- `qos` — Displays all policy-maps of `qos` type.
- `queuing` — Displays all policy-maps configured of `queuing` type.

- `network-qos` — Displays all policy-maps configured of `network-qos` type.
- `control-plane` — Displays all policy-maps of `control-plane` type.
- *`policy-map-name`* — Displays the QoS policy-map name details.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show policy-map
 Service-policy(qos) input: p1
   Class-map (qos): c1
     set qos-group 1
 Service-policy(qos) input: p2
   Class-map (qos): c2
     set qos-group 2
```

Supported Releases 10.2.0E or later

show qos control-plane

Displays the QoS configuration applied to the control-plane.

Syntax `show qos control-plane`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information Monitors statistics for the control-plane and troubleshoots CoPP.

Example

```
OS10# show qos control-plane
     Service-policy (Input): p1
```

Supported Releases 10.2.0E or later

show qos egress bufffers interface

Displays egress buffer configurations.

Syntax `show qos egress buffers interface [interface node/slot/port[:subport]]`

Parameters

- *`interface`* — (Optional) Enter the interface type.
- *`node/slot/port[:subport]`* — (Optional) Enter the port information.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show qos egress buffers interface ethernet 1/1/1
Interface : ethernet1/1/1
Speed : 0
queue-number     pool-type     rsvd-buf-size     threshold-mode     threshold-value
--------------------------------------------------------------------------------------
```

```
0               lossy           1792            dynamic         8
1               lossy           1792            dynamic         8
2               lossy           1792            dynamic         8
3               lossy           1792            dynamic         8
4               lossless        0               dynamic         10
5               lossy           1792            dynamic         8
6               lossy           1792            dynamic         8
7               lossy           1792            dynamic         8
OS10#
```

Supported Releases 10.3.0E or later

show qos egress buffer-statistics-tracking

Displays egress queue-level peak buffer usage count in bytes for queues on a given interface.

Syntax

```
show qos egress buffer-statistics-tracking interface ethernet [node/slot/port]
[[mcast | ucast] queue {all | [0-7]}] [detail]
```

Parameters

- *node/slot/port*—Enter the port information.
- `[[mcast | ucast] queue {all | [0-7]}]`—Enter the `mcast` or `ucast` keyword to view the egress queue peak buffer utilization for multicast or unicast queues respectively. Enter the `all` keyword to specify all queues, or enter the queue number.
- `detail`—Displays per MMU instance-level statistics in platforms with multiple MMU instances.

Default Not applicable

Command Mode EXEC

Usage Information None

Example

```
OS10# show qos egress buffer-statistics-tracking interface
ethernet 1/1/1 mcast queue 0
```

Supported Releases 10.4.3.0 or later

show qos egress buffer-stats interface

Displays the buffers statistics for the egress interface.

Syntax

```
show qos egress buffer-stats interface [interface node/slot/port[:subport]]
```

Parameters

- *interface* — (Optional) Enter the interface type.
- *node/slot/port[:subport]* — (Optional) Enter the port information.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show qos egress buffer-stats interface ethernet 1/1/1
Interface : ethernet1/1/1
Speed : 0
Queue    TX             TX                  Used Total        Used shared
         pckts          bytes               buffers           buffers
---------------------------------------------------------------------------
0         0              0                   0                 0
```

```
1          0              0              0              0
2          0              0              0              0
3          0              0              0              0
4          0              0              0              0
5          0              0              0              0
6          0              0              0              0
7          0              0              0              0
```

Supported Releases 10.3.0E or later

show qos headroom-pool buffer-statistics-tracking

Displays headroom-pool level peak buffer usage count in bytes.

Syntax
```
show qos headroom-pool buffer-statistics-tracking [detail]
```

Parameters detail—Displays headroom-pool statistics per memory management unit (MMU) instance in platforms with multiple MMU instances.

Default Not configured

Command Mode EXEC

Usage Information Supported platforms include Z9100-ON series, Z9200-ON series, S5200-ON series, and MX9116n.

Example
```
OS10# show qos headroom-pool buffer-statistics-tracking
```

Supported Releases 10.4.3.0 or later

show qos ingress buffers interface

Displays interface buffer configurations.

Syntax
```
show qos ingress buffers interface [interface node/slot/port[:subport]]
```

Parameters

- *interface* — (Optional) Enter the interface type.
- *node/slot/port[:subport]* — (Optional) Enter the port information.

Default Not configured

Command Mode EXEC

Usage Information None

Example
```
OS10# show qos ingress buffers interface
Interface : ethernet1/1/1
Speed : 0
Priority-grp  Reserved      Shared-buffer  Shared-buffer  XOFF       XON
    no        buffer-size   mode           threshold      threshold  threshold
--------------------------------------------------------------------------------
0               -             -                -          -            -
1               -             -                -          -                -
2               -             -                -          -                -
3               -             -                -          -                -
4             145152          -                -        98304          89088
5               -             -                -          -                -
6               -             -                -          -                -
7               -             -                -          -                -
```

Supported Releases 10.3.0E or later

show qos ingress buffer-statistics-tracking

Displays ingress priority group-level peak buffer usage count in bytes for the given priority group on a given interface.

Syntax

```
show qos ingress buffer-statistics-tracking interface ethernet [node/slot/port]
[priority-group {0-7}] [detail]
```

Parameters

- *node/slot/port*—Enter the port information.
- `[priority-group {0-7}]`—Enter the `priority-group` keyword, followed by the group number.
- `detail`—Displays per MMU instance-level statistics in platforms with multiple MMU instances.

Default Not applicable

Command Mode EXEC

Usage Information When BST is enabled, if you make any configuration changes that affect the priority group or priority mapping configuration, such as removal of class map, addition of class map to policy map (nqos), and so on, be sure to clear the buffer statistics using the `clear qos statistics type buffer-statistics-tracking` command to view the actual peak buffer utilization for the current configuration.

Example

```
OS10# show qos ingress buffer-statistics-tracking interface
ethernet 1/1/1 priority-group 0
```

Supported Releases 10.4.3.0 or later

show qos ingress buffer-stats interface

Displays the buffers statistics for the ingress interface.

Syntax

```
show qos ingress buffer-stats interface [interface node/slot/port[:subport]]
```

Parameters

- *interface* — (Optional) Enter the interface type.
- *node/slot/port[:subport]* — (Optional) Enter the port information.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show qos ingress buffer-stats interface ethernet 1/1/1
Interface : ethernet1/1/1
Speed : 0
Priority Used Total        Used HDRM
 Group   buffers           buffers
-------------------------------------------------
0         0                 0
1         0                 0
2         0                 0
3         0                 0
4         0                 0
5         0                 0
```

```
6           0                   0
7           0                   0
```

Supported Releases	10.3.0E or later

show qos-rate-adjust

Displays the status of the rate adjust limit for policing and shaping.

Syntax	`show qos-rate-adjust`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	None

Example	`OS10# show qos-rate-adjust`
	
	`QoS Rate adjust configured for Policer and Shaper (in bytes) : 10`

Supported Releases	10.4.3.0 or later

show qos service-pool buffer-statistics-tracking

Displays service-pool level peak buffer usage count in bytes.

Syntax	`show qos service-pool buffer-statistics-tracking [detail]`
Parameters	detail—Displays service-pool level peak buffer utilization per memory management unit (MMU) instance in platforms with multiple MMU instances.
Default	Not configured
Command Mode	EXEC
Usage Information	None
Example	`OS10# show qos service-pool buffer-statistics-tracking`
Supported Releases	10.4.3.0 or later

show qos system

Displays the QoS configuration applied to the system.

Syntax	`show qos system`
Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	View and verify system-level service-policy configuration information.

Example

```
OS10# show qos system

ETS Mode : off
ECN Mode : off □ shows whether the ECN is enabled globally or not
Service-policy (Input) (qos) : policy1
Service-policy (Output)(queuing) : policy2
```

Supported Releases 10.4.1.0 or later

show qos system buffers

Displays the system buffer configurations and utilization.

Syntax `show qos system {ingress | egress} buffers`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show qos system ingress buffer
All values are in kb
Total buffers                                - 12187
    Total lossless buffers                   - 0
      Maximum lossless buffers               - 5512
      Total shared lossless buffers          - 0
      Total used shared lossless buffers     -
    Total lossy buffers                      - 11567
      Total shared lossy buffers             - 11192
      Total used shared lossy buffers        - 0
```

The following command is supported on Z9100-ON and Z9264F-ON platforms:

```
OS10# show qos system ingress buffer detail
All values are in kb
Total buffers                                - 43008
    Total lossless buffers                   - 0
      Maximum lossless buffers               - 23312
      Total shared lossless buffers          - 0
      Total used shared lossless buffers     -
    Total lossy buffers                      - 42388
      Total shared lossy buffers             - 39974
      Total used shared lossy buffers        - 0
      MMU 0
        Total lossy buffers                   - 10597
          Total shared lossy buffers          - 10012
          Total used shared lossy buffers     - 0
      MMU 1
        Total lossy buffers                   - 10597
          Total shared lossy buffers          - 10012
          Total used shared lossy buffers     - 0
      MMU 2
        Total lossy buffers                   - 10597
          Total shared lossy buffers          - 9993
          Total used shared lossy buffers     - 0
      MMU 3
        Total lossy buffers                   - 10597
          Total shared lossy buffers          - 9993
          Total used shared lossy buffers     - 0
```

```
OS10# show qos system egress buffer
All values are in kb
```

```
Total buffers                                - 12187
    Total lossless buffers                   - 0
      Total shared lossless buffers          - 0
      Total used shared lossless buffers     -
    Total lossy buffers                      - 11567
      Total shared lossy buffers             - 9812
      Total used shared lossy buffers        - 0
    Total CPU buffers                      - 620
      Total shared CPU buffers             - 558
      Total used shared CPU buffers          - 0
```

The following command is supported on Z9100-ON and Z9264F-ON platforms:

```
OS10# show qos system egress buffer detail
All values are in kb
Total buffers                                - 43008
    Total lossless buffers                   - 0
      Total shared lossless buffers          - 0
      Total used shared lossless buffers     -
    Total lossy buffers                      - 42388
      Total shared lossy buffers             - 33938
      Total used shared lossy buffers        - 0
      MMU 0
        Total lossy buffers                      - 10597
          Total shared lossy buffers             - 8484
          Total used shared lossy buffers        - 0
      MMU 1
        Total lossy buffers                      - 10597
          Total shared lossy buffers             - 8484
          Total used shared lossy buffers        - 0
      MMU 2
        Total lossy buffers                      - 10597
          Total shared lossy buffers             - 8484
          Total used shared lossy buffers        - 0
      MMU 3
        Total lossy buffers                      - 10597
          Total shared lossy buffers             - 8484
          Total used shared lossy buffers        - 0
```

Supported Releases 10.3.0E or later

show qos maps

Displays the active system trust map.

Syntax `show qos maps type {tc-queue | trust-map-dot1p | trust-map dscp} trust-map-name`

Parameters

- `dot1p` — Enter to view the dot1p trust map.
- `dscp` — Enter to view the DSCP trust map.
- `tc-queue`—Enter to view the traffic class to queue map.
- *`trust-map`* — Enter the name of the trust map.

Default Not configured

Command Mode EXEC

Usage Information None

Example (dot1p)

```
OS10# show qos maps type tc-queue queue-map1
Traffic-Class to Queue Map: queue-map1
Queue        Traffic-Class
--------------------------
```

```
1               5
2               6
3               7
OS10# show qos maps type trust-map-dot1p dot1p-trustmap1
DOT1P Priority to Traffic-Class Map : dot1p-trustmap1
Traffic-Class    DOT1P Priority
--------------------------------
0                   2
1                   3
2                   4
3                   5
4                   6
5                   7
6                   1
OS10# show qos maps type trust-map-dscp dscp-trustmap1
DSCP Priority to Traffic-Class Map  : dscp-trustmap1
Traffic-Class    DSCP Priority
--------------------------------
0                   8-15
2                   16-23
1                   0-7
OS10# show qos maps
Traffic-Class to Queue Map: queue-map1
Queue        Traffic-Class
--------------------------
1               5
2               6
3               7
DOT1P Priority to Traffic-Class Map : map1
Traffic-Class    DOT1P Priority
--------------------------------
DOT1P Priority to Traffic-Class Map : dot1p-trustmap1
Traffic-Class    DOT1P Priority
--------------------------------
0                   2
1                   3
2                   4
3                   5
4                   6
5                   7
6                   1
DSCP Priority to Traffic-Class Map  : dscp-trustmap1
Traffic-Class    DSCP Priority
--------------------------------
0                   8-15
2                   16-23
1                   0-7
Default Dot1p Priority to Traffic-Class Map
Traffic-Class    DOT1P Priority
--------------------------------
0                   1
1                   0
2                   2
3                   3
4                   4
5                   5
6                   6
7                   7
Default Dscp Priority to Traffic-Class Map
Traffic-Class    DSCP Priority
--------------------------------
0                   0-7
1                   8-15
2                   16-23
3                   24-31
4                   32-39
5                   40-47
6                   48-55
7                   56-63
```

```
Default Traffic-Class to Queue Map
Traffic-Class    Queue number
-------------------------------
0                   0
1                   1
2                   2
3                   3
4                   4
5                   5
6                   6
7                   7
OS10#
```

Example (dscp)

```
OS10# show qos trust-map dscp new-dscp-map

new-dscp-map
qos-group   Dscp
   Id
-------------------
0            0-7
1            8-15
2           16-23
3           24-31
4           32-39
5           40-47
6           48-55
7           56-63
```

Supported Releases 10.3.0E or later

show qos wred-profile

Displays the details of WRED profile configuration.

Syntax `show qos wred—profile [wred-profile-name]`

Parameters *wred-profile-name* — (Optional) Enter the Ethernet interface information.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show qos wred-profile
Profile Name         |        Green          |          Yellow      |      Red    |
MIN   MAX   DROP-RATE | MIN    MAX   DROP-RATE |MIN    MAX   DROP-RATE |  WEIGHT| ECN|
 KB    KB           % |  KB     KB           % | KB     KB           % |        |    |
--------------------|----------------------|----------------------|--------|----|
```

Example (S4200) — When ECN is enabled globally.

```
OS10#show qos wred-profile wred_prof1
Wred-profile-name gmin-thd gmax-thd gmax-drop-rate ymin-thd ymax-thd ymax-drop-rate rmin-thd rmax-thd rmax-dr
rate
-------------------------------------------------------------------------------------------------------------------
wred_prof1 0 0 0 1 10 40 0 0 0

S4200 o/p



OS10# show qos wred-profile

Profile Name |       Green           |      Yellow          |              Red                  |
-------------|-----------------------|----------------------|-----------------------------------|
             |  MIN  MAX DROP-RATE   |  MIN  MAX DROP-RATE  | MIN  MAX DROP-RATE | WEIGHT |  ECN|
             |   KB   KB          %  |   KB   KB          % |  KB   KB          % |        |     |
-------------|-----------------------|----------------------|--------------------|--------|-----|

profile1     |   10  100       100   |                      |                    |        |  Off|
```

```
-------------|----------------------|--------------------|--------------------|--------|-----|
profile2     |                      |                    |                    |        |  On|
             |----------------------|--------------------|--------------------|--------|-----|
Color Blind ECN Thd| 100 1000   100  |
-------------|----------------------|--------------------|--------------------|--------|-----|
```

Supported Releases

show queuing statistics

Displays QoS queuing statistics information.

Syntax `show queuing statistics interface ethernet` *node/slot/port[:subport]* [queue *number*]

Parameters

- *node/slot/port[:subport]* — Enter the Ethernet interface information.
- queue *number* — Enter the QoS queue number, from 0 to 7.

Default Not configured

Command Mode EXEC

Usage Information Use this command to view all queuing counters. WRED counters are available only at the port level.

Example

```
OS10# show queuing statistics interface ethernet 1/1/1
Interface ethernet1/1/1 (All queues)
Description Packets  Bytes
Output          0       0
Dropped         0       0
Green Drop      0       0
Yellow Drop     0       0
Red drop        0       0
```

Example (Queue)

```
OS10# show queuing statistics interface ethernet 1/1/1 queue 3
Interface ethernet1/1/1 Queue 3
Description Packets  Bytes
Output          0       0
Dropped         0       0
```

Supported Releases 10.2.0E or later

system qos

Enters SYSTEM-QOS mode to configure system-level service policies.

Syntax `system qos`

Parameters None

Default Not configured

Command Mode CONFIGURATION

Usage Information None

Example

```
OS10(config)# system qos
OS10(config-sys-qos)#
```

Supported Releases 10.2.0E or later

trust-map

Configures trust map on an interface or on a system QoS.

Syntax `trust—map {dot1p | dscp} {default | trust-map-name}`

Parameters

- `dot1p` — Apply dot1p trust map.
- `dscp` — Apply dscp trust map.
- `default` — Apply default dot1p or dscp trust map.
- *`trust-map-name`* — Enter the name of trust map.

Default Disabled

Command Mode INTERFACE

SYSTEM-QoS

Usage Information The `no` version of this command removes the applied trust map from the interface or system QoS.

Example

```
OS10(config)# interface ethernet 1/1/10
OS10(conf-if-eth1/1/10)# trust-map dot1p default
OS10(conf-if-eth1/1/10)# trust-map dot1p d1

OS10(config)# interface ethernet 1/1/2
OS10(conf-if-eth1/1/2)# trust-map dscp default
OS10(conf-if-eth1/1/2)# trust-map dscp d2
```

```
OS10(config)# system qos
OS10(config-sys-qos)# trust-map dot1p default
OS10(config-sys-qos)# trust-map dscp d2
```

Supported Releases 10.4.1.0 or later

trust dot1p-map

Creates a user-defined trust map for dot1p flows.

Syntax `trust dot1p-map map-name`

Parameters *`map-name`* — Enter the name of the dot1p trust map. A maximum of 32 characters.

Default Not configured

Command Mode CONFIGURATION

Usage Information If you enable trust, traffic obeys the dot1p map. `default-dot1p-trust` is a reserved trust-map name. The `no` version of this command returns the value to the default.

Example

```
OS10(config)# trust dot1p-map map1
OS10(config-tmap-dot1p-map)# qos-group 4 dot1p 5
```

Supported Releases 10.3.0E or later

trust dscp-map

Creates a user-defined trust map for DSCP flows.

Syntax

```
trust dscp-map map-name
```

Parameters

map-name — Enter the name of the DSCP trust map. A maximum of 32 characters.

Default

Not configured

Command Mode

CONFIGURATION

Usage Information

If you enable trust, traffic obeys this trust map. `default-dscp-trust` is a reserved trust-map name. The `no` version of this command returns the value to the default.

Example

```
OS10(config)# trust dscp-map dscp-trust1
```

Supported Releases

10.3.0E or later

qos-map traffic-class

Creates a user-defined trust map for queue mapping. In S5148F-ON, apply the traffic class only on the egress traffic.

Syntax

```
qos-map traffic-class map-name
```

Parameters

map-name — Enter the name of the queue trust map. A maximum of 32 characters.

Default

Not configured

Command Mode

CONFIGURATION

Usage Information

If applied on the interface or system level, the traffic class routes all traffic to the mapped queue. The `no` version of this command returns the value to the default.

Example

```
OS10(config)# qos-map traffic-class queue-map1
OS10(config-qos-map)# queue 1 qos-group 5
OS10(config-qos-map)# queue 2 qos-group 6
OS10(config-qos-map)# queue 3 qos-group 7
OS10(config-qos-map)#
```

Supported Releases

10.3.0E or later

trust-map

Applies a dot1p or DSCP traffic class to a queue trust map.

Syntax

```
trust {dot1p | dscp} {default | trust-map-name}
```

Parameters

- `dot1p`— Applies a dot1p trust map.
- `dscp`—Applies a dscp trust map.
- `default`— Applies a default trust map.

Default

Disabled

Command Mode	SYSTEM-QOS INTERFACE
Usage Information	Use the `show qos maps type [tc-queue | trust-map-dot1p | trust-map-dscp] [string]` command to view the current trust mapping. You must change the trust map only during no traffic flow. Verify the correct policy maps are applied. The `no` version of this command returns the value to the default.
Example	`OS10(config)# interface ethernet 1/1/1` `OS10(conf-if-eth1/1/1)# trust-map dscp dscp-trustmap1`
Supported Releases	10.4.1.0 or later

wred

Configures a weighted random early detection (WRED) profile.

Syntax	`wred` *`wred-profile-name`*
Parameters	*`wred-profile-name`* — Enter a name for the WRED profile.
Default	Not configured
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command removes the WRED profile.
Example	`OS10(config)# wred test_wred` `OS10(config-wred)#`
Supported Releases	10.4.0E(R1) or later

13

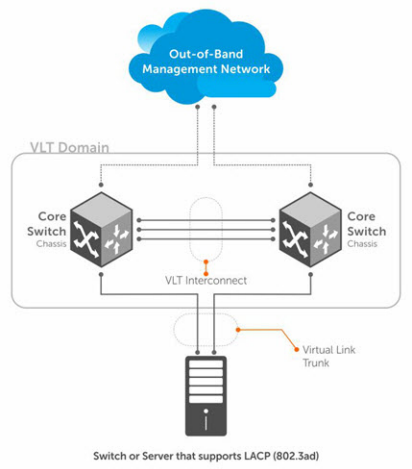
Virtual Link Trunking

Virtual Link Trunking (VLT) is a Layer 2 (L2) aggregate protocol between end devices such as servers connected to different network devices. VLT reduces the role of Spanning Tree Protocols (STPs) by allowing link aggregation group (LAG) terminations on two separate distributions or core switches.

VLT:

- Allows a single device to use a LAG across two upstream devices
- Provides a loop-free topology
- Eliminates STP-blocked ports
- Optimizes using all available uplink bandwidth
- Guarantees fast convergence if either a link or device fails
- Enhances optimized forwarding with Virtual Router Redundancy Protocol (VRRP)
- Provides link-level resiliency
- Assures high availability

VLT provides L2 multipathing, creating redundancy through increased bandwidth, enabling multiple parallel paths between nodes and load-balancing traffic where alternative paths exist.



VLT presents a single logical L2 domain from the perspective of attached devices that have a virtual link trunk terminating on separate nodes in the VLT domain. The two VLT nodes are independent Layer2/ Layer3 (L2/L3) switches for devices in the upstream network. L2/L3 control plane protocols and system management features function normally in VLT mode.

VLT configurations must be identical on both sides of a trunk. External switches or servers with LACP see the VLT switches as a single virtual switch.

VLT physical ports	802.1p, 802.1q, LLDP, flow control, port monitoring, and jumbo frames are supported on VLT physical ports.
System management protocols	All system management protocols are supported on VLT ports — SNMP, RMON, AAA, ACL, DNS, FTP, SSH, syslog, NTP, RADIUS, SCP, and LLDP.

L3 VLAN connectivity	Enable L3 VLAN connectivity, VLANs assigned with an IP address, on VLT peers by configuring a VLAN interface for the same VLAN on both devices.
Optimized forwarding with VRRP	To ensure the same behavior on both sides of the VLT nodes, VRRP requires state information coordination. VRRP Active-Active mode optimizes L3 forwarding over VLT. By default, VRRP Active-Active mode is enabled on all the VLAN interfaces. VRRP Active-Active mode enables each peer to locally forward L3 packets, resulting in reduced traffic flow between peers over the VLTi link.
Spanning-Tree Protocol	VLT ports support RSTP, RPVST+, and MSTP.
Multicast	IGMP snooping and MLD snooping are supported on VLT ports.

NOTE: 802.1x and DHCP snooping are not supported on VLT ports.

Terminology

VLT domain	The domain includes VLT peer devices, VLT interconnect, and all port-channels in the VLT connected to the attached devices. It is also the configuration mode that you must use to assign VLT global parameters.
VLT interconnect (VLTi)	The link between VLT peer switches used to synchronize operating states.
VLT peer device	A pair of devices connected using a dedicated port-channel — the VLTi. You must configure VLT peers separately.
Discovery interface	Port interfaces on VLT peers in the VLT interconnect (VLTi) link.
VLT MAC address	(Optional) Unique MAC address that you assign to the VLT domain. A VLT MAC address is the common address all VLT peers use. If you do not configure a VLT MAC address, the MAC address of the primary peer is used as the VLT MAC address across all peers.
VLT node priority	The priority based on which the primary and secondary VLT nodes are determined. If priority is not configured, the VLT node with the lowest MAC address is elected as the primary VLT node.
VLT port-channel	A combined port-channel between an attached device and VLT peer switches.
VLT port-channel ID	Groups port-channel interfaces on VLT peers into a single virtual-link trunk connected to an attached device. Assign the same port-channel ID to interfaces on different peers that you bundle together.
Orphan ports	Ports that are connected to VLT domain, but not part of the VLT-LAG.

VLT domain

A VLT domain includes the VLT peer devices, VLTi, and all VLT port-channels that connect to the attached devices. It is also the configuration mode that you must use to assign VLT global parameters.

- Each VLT domain must have a unique MAC address that you create or that VLT creates automatically.
- VLAN ID 4094 is reserved as an internal control VLAN for the VLT domain.
- ARP, IPv6 neighbors, and MAC tables synchronize between the VLT peer nodes.
- VLT peer devices operate as a separate node with independent control and data planes for devices that attach to non-VLT ports.
- One node in the VLT domain takes a primary role and the other node takes the secondary role. In a VLT domain with two nodes, the VLT assigns the primary node role to the node with the lowest MAC address by default. You can override the default primary election mechanism by assigning priorities to each node using the `primary-priority` command.
- If the primary peer fails, the secondary peer (with the higher priority) takes the primary role. If the primary peer (with the lower priority) later comes back online, it is assigned the secondary role (there is no preemption).
- In a VLT domain, the peer network devices must run the same OS10 software version.
- Configure the same VLT domain ID on peer devices. If a VLT domain ID mismatch occurs on VLT peers, the VLTi does not activate.
- In a VLT domain, VLT peers support connections to network devices that connect to only one peer.

VLT interconnect

A VLT interconnect (VLTi) synchronizes states between VLT peers. OS10 automatically adds VLTi ports to VLANs spanned across VLT peers and does not add VLTi ports to VLANs configured on only one peer.

- VLAN ID 4094 is reserved as an internal control VLAN for the VLT domain, and it is not user configurable.
- The VLTi synchronizes L2 and L3 control-plane information across the two nodes. The VLTi is used for data traffic only when there is a link failure that requires VLTi to reach the final destination.
- Traffic with an unknown destination MAC address, multicast, or broadcast traffic can cause flooding across the VLTi.
- MAC, ARP, IPv6 neighbors that are learned over VLANs on VLT peer nodes synchronize using VLTi.
- LLDP, flow control, port monitoring, and jumbo frame features are supported on a VLTi.

Graceful LACP with VLT

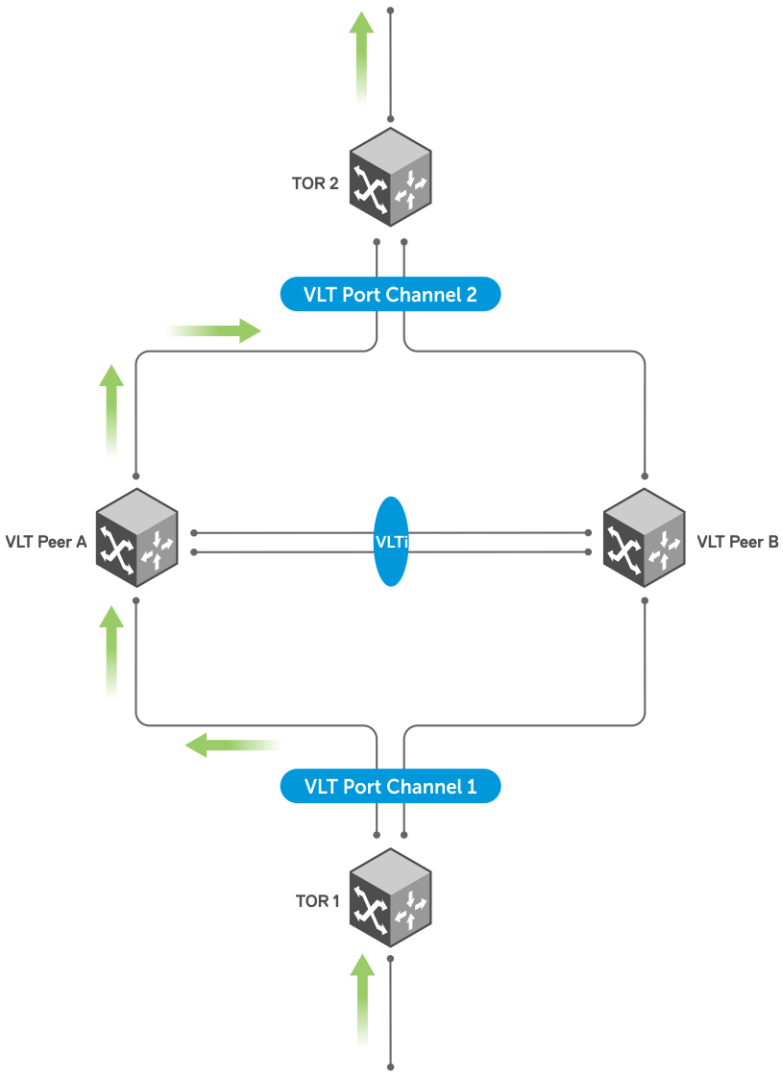
When a VLT node is reloaded, all its interfaces including VLT port-channel interfaces go down. Top-of-rack (ToR) devices connected at the other end of the VLT port-channel interfaces take a considerable amount of time to detect the interface status change and switch the traffic towards the other active VLT node. This causes traffic loss for a long time interval. Using LACP PDUs, the graceful LACP feature allows VLT nodes to inform ToR devices ahead of taking down the member ports of its VLT port-channel interfaces. This enables the ToR devices to switch the traffic to the other active VLT node.

Graceful LACP is supported in these scenarios:

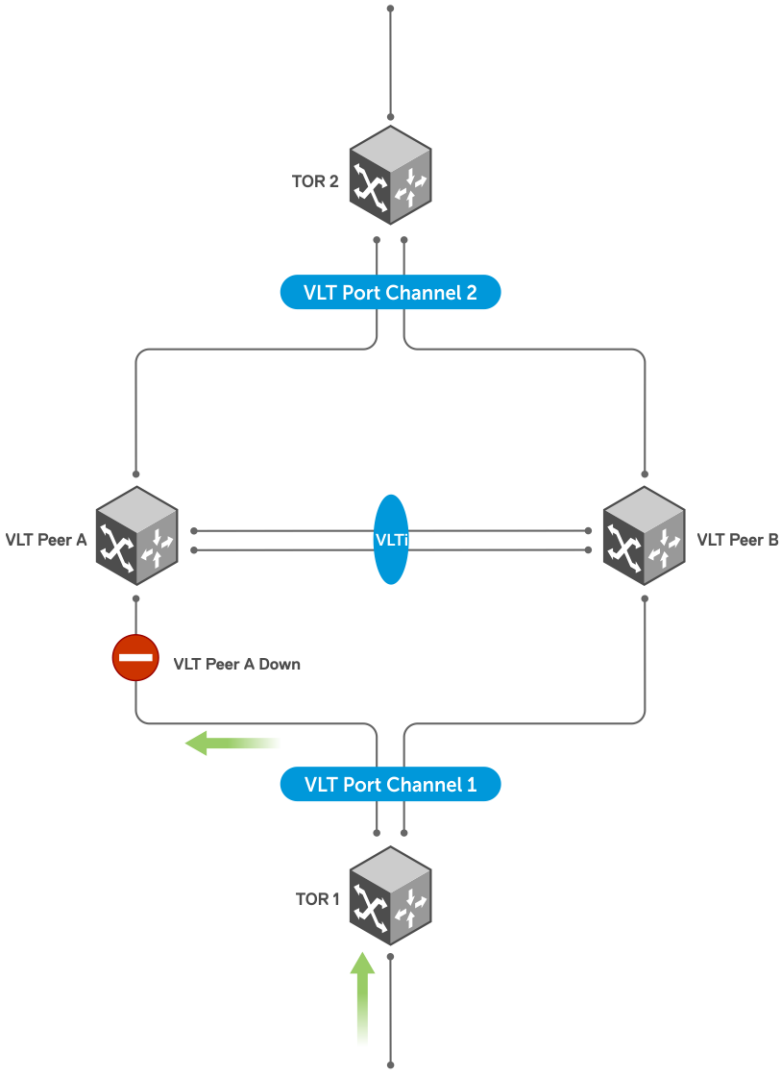
- When a VLT node is reloaded.
- When the secondary VLT node detects that the VLTi link is down but the heartbeat is functional.

Graceful LACP is enabled by default and you cannot disable it.

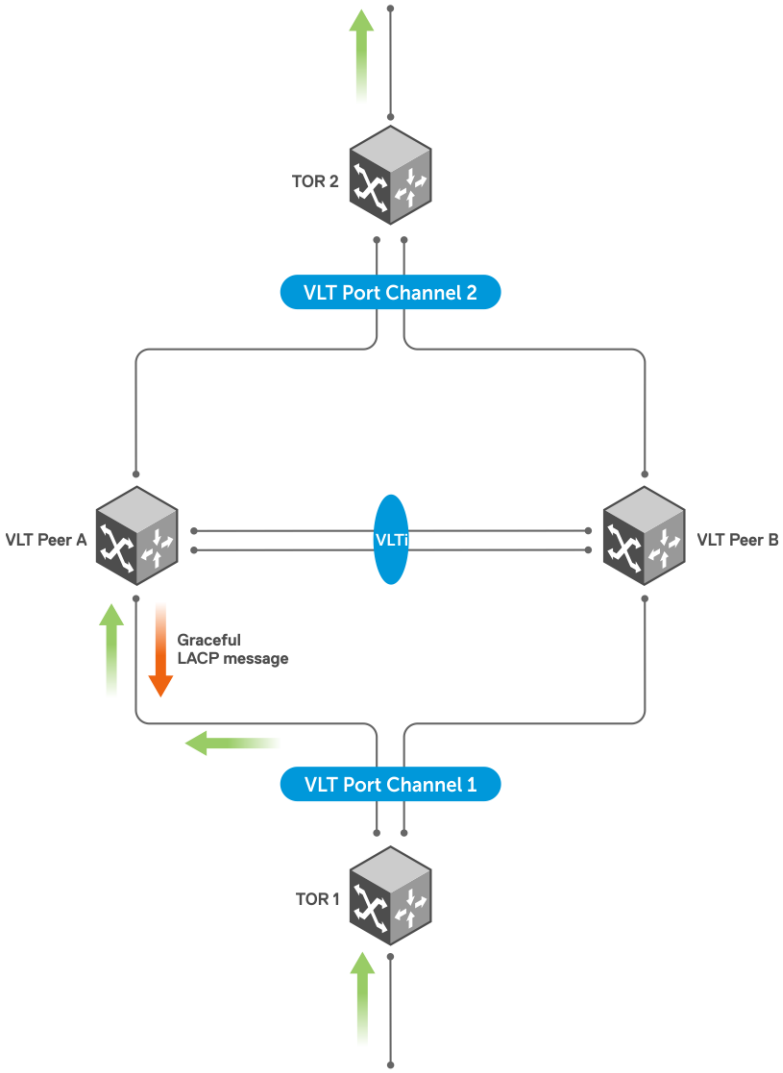
The following shows the normal behavior of a VLT setup where data flows through the optimal path to its destination.



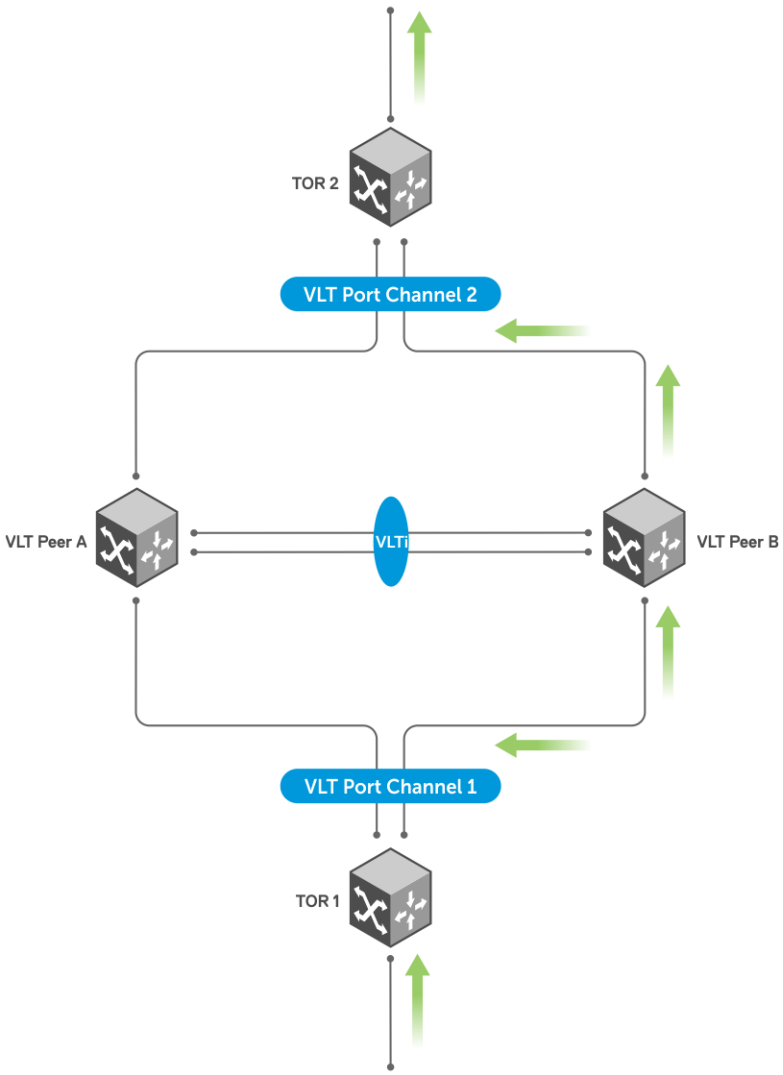
The following shows a scenario where the forwarding device, VLT Peer A, is being reloaded or going down. Until LACP convergence happens, the ToR device continues to forward traffic to VLT Peer A resulting in traffic loss for a longer time interval.



With graceful LACP, VLT Peer A sends graceful LACP PDUs out to all VLT member ports, as shown:



These PDUs notify ToR 1 to immediately direct the traffic to VLT Peer B and minimizing traffic loss.



Configure VLT

Verify that both VLT peer devices are running the same OS version. For VRRP operation, configure VRRP groups and L3 routing on each VLT peer.

Configure the following settings on each VLT peer device separately.

1 (Optional) To prevent loops in VLT domain, enable the STP globally using the `spanning-tree mode {rstp | rapid-pvst | mst}` command.

2 Create a VLT domain by configuring the same domain ID on each peer using the `vlt-domain` command.

3 (Optional) To override the default VLT primary election mechanism based on the system MAC addresses of the VLT nodes, configure a VLT node priority for each of the VLT nodes using the `primary-priority` command. Enter a lower priority value for the desired primary VLT peer and a higher priority value for the desired secondary VLT peer.

4 Configure the VLTi interfaces on each peer using the `discovery-interface` command. After you configure both sides of the VLTi, the primary and secondary roles in the VLT domain are automatically assigned if primary priority is not configured.

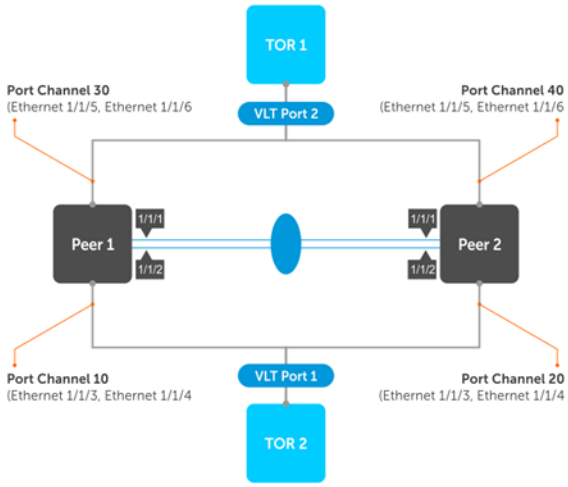
5 (Optional) Manually reconfigure the default VLT MAC address. Configure the VLT MAC address in both VLT peers.

6 (Optional) Configure a time interval to delay bringing up VLT ports after reload or when VLTi come up after failure.

7 Configure the VLT backup link the heartbeat use with the `backup destination {`*`ip-address`*` | ipv6 `*`ipv6-address`*` } [vrf management] [interval `*`interval-time`*`]` command.

8 Configure VLT port-channels between VLT peers and an attached device using the `vlt-port-channel` command. Assign the same VLT port-channel ID from 1 to 1024 to interfaces on different peers that you bundle together. The peer interfaces appear as a single VLT LAG to downstream devices.

9 Connect peer devices in a VLT domain to an attached access device or server.



MSTP configuration

When you enable Multiple Spanning Tree Protocol (MSTP) on VLT nodes, configure both VLT peer nodes in the same MST region to avoid network loops. Ensure that the VLAN-to-instance mappings, region name, and revision ID are same on both VLT peer nodes.

To configure MSTP over VLT, follow these steps on both VLT peer nodes:

1 Enable MSTP.
CONFIGURATION mode

```
spanning-tree mode mst
```

2 Enter MST configuration mode.
CONFIGURATION mode

```
spanning tree mst configuration
```

3 Create an MST instance and add multiple VLANs as required.
MULTIPLE-SPANNING-TREE

```
instance instance-number vlan from-vlan-id – to-vlan-id
```

4 Configure the MST revision number, from 0 to 65535.
MULTIPLE-SPANNING-TREE

```
revision revision-number
```

5 Configure the MST region name.
MULTIPLE-SPANNING-TREE

name *name-string*

The following example shows that both VLT nodes are configured with the same MST VLAN-to-instance mapping..

VLT Peer 1 confoguration

```
OS10(config)# spanning-tree mode mst
```

```
OS10(config)# spanning-tree mst configuration
OS10(conf-mst)# instance 1 vlan 2-10
```

VLT Peer 2 confoguration

```
OS10(config)# spanning-tree mode mst
OS10(config)# spanning-tree mst configuration
OS10(conf-mst)# instance 1 vlan 2-10
```

The following example shows MSTP information on VLT:

```
OS10# show spanning-tree virtual-interface
VFP(VirtualFabricPort) of MSTI 0 is Designated Forwarding
Edge port: No (default)
Link type: point-to-point (auto)
Boundary: Yes, Bpdu-filter: Disable, Bpdu-Guard: Disable, Shutdown-on-Bpdu-Guard-violation: No
Root-Guard: Disable, Loop-Guard: Disable
Bpdus (MRecords) Sent: 387, Received: 16
Interface                                                                    Designated
Name                PortID    Prio      Cost      Sts         Cost      Bridge ID                 PortID
-------------------------------------------------------------------------------------------------------
VFP(VirtualFabricPort) 0.1     0         1         FWD         0         32768    3417.ebf2.a8c4  0.1

VLT-LAG -1(vlt-portid-1) of MSTI 0 is in Designated Forwarding
Edge port: No (default)
Link type: point-to-point (auto)
Boundary: No, Bpdu-filter: Disable, Bpdu-Guard: Disable,
Shutdown-on-Bpdu-Guard-violation: No
Root-Guard: Disable, Loop-Guard: Disable
Bpdus (MRecords) Sent: 1234, Received: 123

Virtual
Interface                                                          Designated
Name              PortID    Prio      Cost      Sts         Cost    Bridge ID                 PortID
-------------------------------------------------------------------------------------------------------

VLT-LAG -1(vlt-portid1)         128.2001  128       2000000   FWD         0       32768    90b1.1cf4.a523
128.2001
```

The following example shows MSTP information on VLT in detail:

```
OS10# show spanning-tree virtual-interface detail
Port 1 (VFP(VirtualFabricPort)) of MSTI 0 is designated Forwarding
Port path cost 0, Port priority 128, Port Identifier 128.1
Designated root priority: 32768, address: 34:17:44:55:66:7f
Designated bridge priority: 32768, address: 90:b1:1c:f4:a5:23
Designated port ID: 128.1, designated path cost: 0
Number of transitions to forwarding state: 1
Edge port: No (default)
Link Type: Point-to-Point
BPDU Sent: 2714, Received: 1234

Port 2001 (VLT-LAG -1(vlt-portid-1)) of MSTI 0 is designated Forwarding
Port path cost 200000, Port priority 128, Port Identifier 128.2001
Designated root priority: 32768, address: 34:17:44:55:66:7f
Designated bridge priority: 32768, address: 90:b1:1c:f4:a5:23
Designated port ID: 128.2001, designated path cost: 0
Number of transitions to forwarding state: 1
Edge port: No (default)
Link Type: Point-to-Point
BPDU Sent: 2714, Received: 1234
```

RSTP configuration

RSTP prevents loops during the VLT startup phase. If required, configure RSTP in the network, before you configure VLT on peer switches.

- Enable RSTP on each peer node in CONFIGURATION mode.

  ```
  spanning-tree mode rstp
  ```

Configure RSTP — peer 1

```
OS10(config)# spanning-tree mode rstp
```

Configure RSTP — peer 2

```
OS10(config)# spanning-tree mode rstp
```

View VLT-specific STP information

```
OS10# show spanning-tree virtual-interface
VFP(VirtualFabricPort) of RSTP 1 is Designated Forwarding
Edge port: No (default)
Link type: point-to-point (auto)
Boundary: No, Bpdu-filter: Disable, Bpdu-Guard: Disable, Shutdown-on-Bpdu-Guard-violation: No
Root-Guard: Disable, Loop-Guard: Disable
Bpdus (MRecords) Sent: 11, Received: 7
Interface                                                                  Designated
Name                 PortID   Prio   Cost  Sts   Cost  Bridge ID          PortID
---------------------------------------------------------------------------------
VFP(VirtualFabricPort) 0.1     0     1     FWD    0    32768 0078.7614.6062  0.1
```

View STP virtual interface detail

```
OS10# show spanning-tree virtual-interface detail
Port 1 (VFP(VirtualFabricPort)) of RSTP 1 is designated Forwarding
Port path cost 1, Port priority 0, Port Identifier 0.1
Designated root priority: 32768, address: 00:78:76:14:60:62
Designated bridge priority: 32768, address: 00:78:76:14:60:62
Designated port ID: 0.1, designated path cost: 0
Number of transitions to forwarding state: 1
Edge port: No (default)
Link Type: Point-to-Point
BPDU Sent: 15, Received: 5
```

RPVST+ configuration

Use RPVST+ for initial loop prevention during the VLT startup phase. If required, configure RPVST+ in the network before you configure VLT on peer switches.

Configure RPVST+ on both the VLT peers. This creates an RPVST+ instance for every VLAN configured in the system. The RPVST+ instances in the primary VLT peer control the VLT LAGs on both the primary and secondary peers.

- Enable RPVST+ on each peer node in CONFIGURATION mode.

  ```
  spanning-tree mode rapid-pvst
  ```

Configure RPVST+ — peer 1

```
OS10(config)# spanning-tree mode rapid-pvst
```

Configure RPVST+ — peer 2

```
OS10(config)# spanning-tree mode rapid-pvst
```

View RPVST+ information on VLT

```
OS10# show spanning-tree virtual-interface

VFP(VirtualFabricPort) of vlan 100 is Designated Blocking
Edge port: No (default)
Link type: point-to-point (auto)
Boundary: No, Bpdu-filter: Disable, Bpdu-Guard: Disable, Shutdown-on-Bpdu-Guard-violation: No
Root-Guard: Disable, Loop-Guard: Disable
Bpdus (MRecords) Sent: 7, Received: 9
Interface                                                                  Designated
```

```
Name                    PortID    Prio      Cost      Sts         Cost      Bridge
ID              PortID
-------------------------------------------------------------------------------------------------
-------------
VFP(VirtualFabricPort)  0.1       0         1         BLK         0         4196
90b1.1cf4.a602  0.1
```

View RPVST+ information on VLT in detail

```
OS10# show spanning-tree virtual-interface detail
Port 1 (VFP(VirtualFabricPort)) of vlan1 is designated Forwarding
Port path cost 1, Port priority 0, Port Identifier 0.1
Designated root priority: 4097, address: 90:b1:1c:f4:a6:02
Designated bridge priority: 4097, address: 90:b1:1c:f4:a6:02
Designated port ID: 0.1, designated path cost: 0
Number of transitions to forwarding state: 1
Edge port: No (default)
Link Type: Point-to-Point
BPDU Sent: 202, Received: 42
Port 1 (VFP(VirtualFabricPort)) of vlan100 is designated Forwarding
Port path cost 1, Port priority 0, Port Identifier 0.1
Designated root priority: 4196, address: 90:b1:1c:f4:a6:02
Designated bridge priority: 4196, address: 90:b1:1c:f4:a6:02
Designated port ID: 0.1, designated path cost: 0
Number of transitions to forwarding state: 1
Edge port: No (default)
Link Type: Point-to-Point
BPDU Sent: 101, Received: 21
```

Create VLT domain

A VLT domain requires an ID number. Configure the same VLT domain ID on both peers. For more information, see VLT domain. The `no vlt-domain` command disables VLT.

1 Configure a VLT domain and enter VLT-DOMAIN mode. Configure the same VLT domain ID on each peer, from 1 to 255.

```
vlt-domain domain-id
```

2 Repeat the steps on the VLT peer to create the VLT domain.

Peer 1

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)#
```

Peer 2

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)#
```

VLTi configuration

Before you configure VLTi on peer interfaces, remove each interface from L2 mode with the `no switchport` command. For more information, see VLT interconnect.

1 Enter the VLT domain ID to enter from CONFIGURATION mode.

```
vlt-domain domain-id
```

2 Configure one or a hyphen-separated range of VLT peer interfaces to become a member of the VLTi in INTERFACE mode.

```
discovery-interface {ethernet node/slot/port[:subport] | ethernet node/slot/port[:subport] -
node/slot/port[:subport]}
```

3 Repeat the steps on the VLT peer.

Peer 1

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# exit
OS10(config)# interface ethernet 1/1/2
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/2)# exit
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# discovery-interface ethernet1/1/1
OS10(conf-vlt-1)# discovery-interface ethernet1/1/2
```

Peer 2

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# no switchport
OS10(conf-if-eth1/1/1)# exit
OS10(config)# interface ethernet 1/1/2
OS10(conf-if-eth1/1/2)# no switchport
OS10(conf-if-eth1/1/2)# exit
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# discovery-interface ethernet1/1/1
OS10(conf-vlt-1)# discovery-interface ethernet1/1/2
```

Configure VLT MAC address

You can manually configure the VLT MAC address.

Configure the VLT MAC address symmetrical in both the VLT peer switches to avoid any unpredictable behavior when any unit is down or when VLTi is reset. If you do not configure a VLT MAC address, the MAC address of the primary peer is used as the VLT MAC address across all peers. Configuring the MAC address manually enables to minimize the time required to synchronize the default MAC address of the VLT domain on both peer devices when one peer switch reboots.

Use the `vlt-mac` *`mac-address`* to configure the MAC address in both the VLT peers.

Example configuration:

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# vlt-mac 00:00:00:00:00:02
```

NOTE: It is recommended to configure the same VLT MAC address manually on both the VLT peer switches.

Delay restore timer

When a VLT node boots up, restoration of VLT port status is deferred for a certain amount of time to enable VLT peers to complete the control data information exchange.

If the peer VLT device was up at the time the VLTi link failed, the system allows a delay in bringing up of VLT ports after reload or peer-link restoration between the VLT peer switches..

When both the VLT peers are up and running, and if VLTi fails, the secondary peer brings down the VLT ports. When the VLTi port comes up, secondary peer does not bring up VLT ports immediately. The VLT ports are brought up only after the VLT port restoration timer, to allow both the VLT peers to sync up the control information with each other.

By default, the system allows 90 seconds. You can use the `delay-restore` *`timer`* command to modify the duration of the timer.

Example:

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# delay-restore 100
```

VLT backup

VLT backup link is an additional link used to check the availability of the peer nodes in the VLT domain.

When VLTi interface goes down, the backup link helps to differentiate the VLTi link failure from peer node failure. If the VLTi link fails, all the VLT nodes exchange node liveliness information through the backup link.

Based on the node liveliness information, the VLT LAG/port is in up state in the primary VLT peer and in down state in the secondary VLT peer. When only the VLTi link fails, but the peer is alive, the secondary VLT peer shuts down the VLT ports. When the node in primary peer fails, the secondary becomes the primary peer.

Configure the VLT backup link using the `backup destination {ip-address | ipv6 ipv6-address} [vrf management] [interval interval-time]`

Example configuration:

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination 10.16.151.110 vrf management interval 30

OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination ipv6 1::1 vrf management interval 30
```

The following examples describe different cases where VLT backup link can be used:

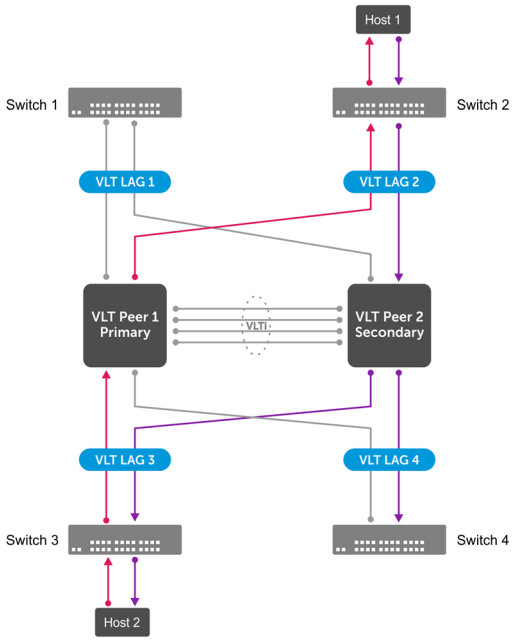
Support for new streams during VLTi failure

When VLTi fails, MAC address learnt after the failure is not synchronized with VLT peers. This leads to continuous flooding of traffic instead of unicast.

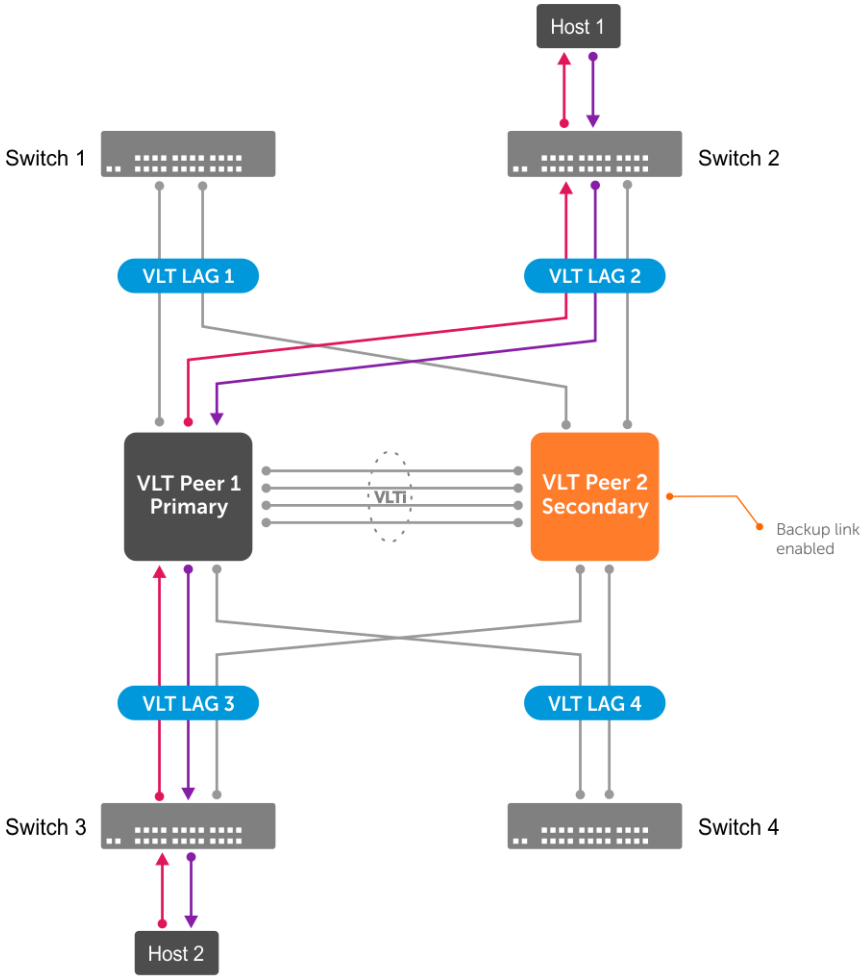
Due to wrong hashing, ARP learning might fail leading to traffic being dropped.

In the following illustration, after VLTi is down VLT peer1 learns MAC address of Host2.

As VLTi link fails, the VLT peer2 is not synched up with the MAC address of Host2 . Due to this, if the traffic from Host1 is hashed to VLT peer2, then the VLT peer2 floods the traffic. Yet, the traffic would not reach Host2 as the VLT port between VLT peer 2 and Swicth 3 is down.



When VLT backup link is enabled, the secondary VLT peer 2 identifies the node liveliness through the backup link. If the primary is up, the secondary peer brings down the VLT LAG ports. Now the traffic from Host1 reaches VLT peer 1 and then reaches the destination, that is Host2. In this case the traffic is unicasted instead of flooding, as shown in the following illustration.

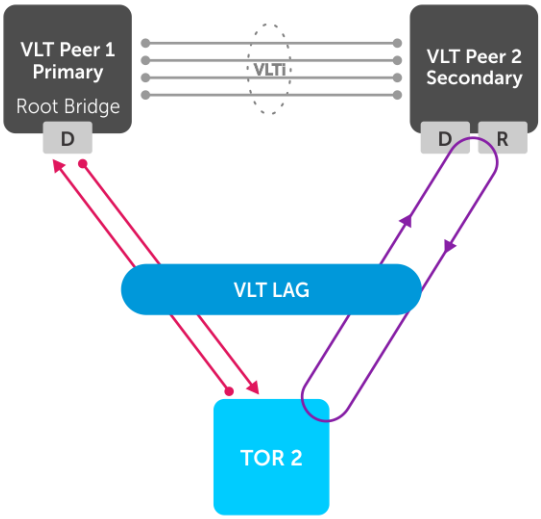


Prevention of loops during VLTi failure

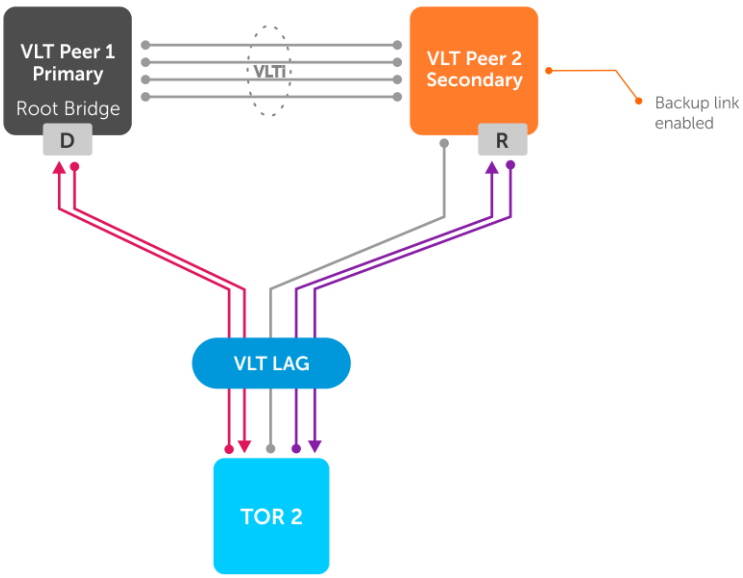
When VLTi is down, STP may fail to detect any loops in the system, which creates data loop in an L2 network.

In the following illustration, STP is running in all the three switches. In the steady state, VLT peer 1 is elected as the root bridge.

When VLTi is down, both the VLT nodes become primary. In this state, VLT peer 2 sends STP BPDU to TOR assuming that TOR sends BPDU to VLT peer 1. Due to this, VLT peer 2 does not receive BPDU on the VLT port, but receives TOR BPDU from orphan port. The STP in VLT peer 2 assumes that there is no loop in the system and opens up both the VLT and the orphan ports. This creates a data loop in the system which brings down the system.



When VLT backup link is enabled, the secondary VLT peer identifies the node liveliness of primary through the backup link. If the primary VLT peer is alive, the secondary VLT peer brings down the VLT LAG ports. In this scenario, the STP opens up the orphan port and there is no loop in the system as shown in the following illustration.



Configure VLT port-channel

A VLT port-channel links an attached device and VLT peer switches, also known as a virtual link trunk. OS10 supports a maximum of 128 VLT LAG port-channels per node.

1 Enter the port-channel ID number on the VLT peer in INTERFACE mode, from 1 to 1024.

```
interface port-channel id-number
```

2 Assign the same ID to a VLT port-channel on each VLT peer. The peers are seen as a single VLT LAG to downstream devices.

```
vlt-port-channel vlt-lag-id
```

3 Repeat the steps on the VLT peer.

Configure VLT LAG — peer 1

```
OS10(config)# interface port-channel 10
OS10(conf-if-po-10)# vlt-port-channel 1
```

Configure VLT LAG — peer 2

```
OS10(config)# interface port-channel 20
OS10(conf-if-po-20)# vlt-port-channel 1
```

VLT unicast routing

VLT unicast routing enables optimized routing where packets destined for the L3 endpoint of the VLT peer are locally routed. IPv4 and IPv6 support VLT unicast routing.

To enable VLT unicast routing, both VLT peers must be in L3 mode. The VLAN configuration must be symmetrical on both peers. You cannot configure the same VLAN as L2 on one node and as L3 on the other node.

1 Enter the VLT domain ID in CONFIGURATION mode, from 1 to 1024.

```
vlt-domain domain-id
```

2 Enable peer-routing in VLT-DOMAIN mode.

```
peer-routing
```

3 Repeat the steps on the VLT peer.

Configure unicast routing — peer 1

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# peer-routing
```

View unicast routing — peer 1

```
do show running-configuration vlt
!
vlt-domain 1
 discovery-interface ethernet1/1/3-1/1/6,1/1/53:1-1/1/53:4,1/1/54:1-1/1/54:4
 peer-routing
```

Configure unicast routing — peer 2

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# peer-routing
```

View unicast routing — peer 2

```
do show running-configuration vlt
!
vlt-domain 1
 discovery-interface ethernet1/1/3-1/1/6,1/1/53:1-1/1/53:4,1/1/54:1-1/1/54:4
 peer-routing
```

VRRP Optimized Forwarding

To enable optimized L3 forwarding over VLT, use VRRP Active-Active mode. By default, VRRP Active-Active mode is enabled on the VLAN interfaces. In this mode, each peer locally forwards L3 traffic, resulting in reduced traffic flow over the VLTi. Configure the same static and dynamic L3 routing on each peer to ensure that L3 reachability and routing tables are the same on both peers.

1 Enable VRRP Active-Active mode in VLAN-INTERFACE mode.

```
vrrp mode active-active
```

2 Configure VRRP on the L3 VLAN that spans both peers.

3 Repeat the steps on the VLT peer.

Configure VRRP active-active mode — peer 1

```
OS10(conf-if-vl-10)# vrrp mode active-active
```

Configure VRRP active-active mode — peer 2

```
OS10(conf-if-vl-10)# vrrp mode active-active
```

View VRRP configuration

```
OS10# show running-configuration interface vlan 10
!
interface vlan10
 no shutdown
 vrrp mode active-active
OS10#
```

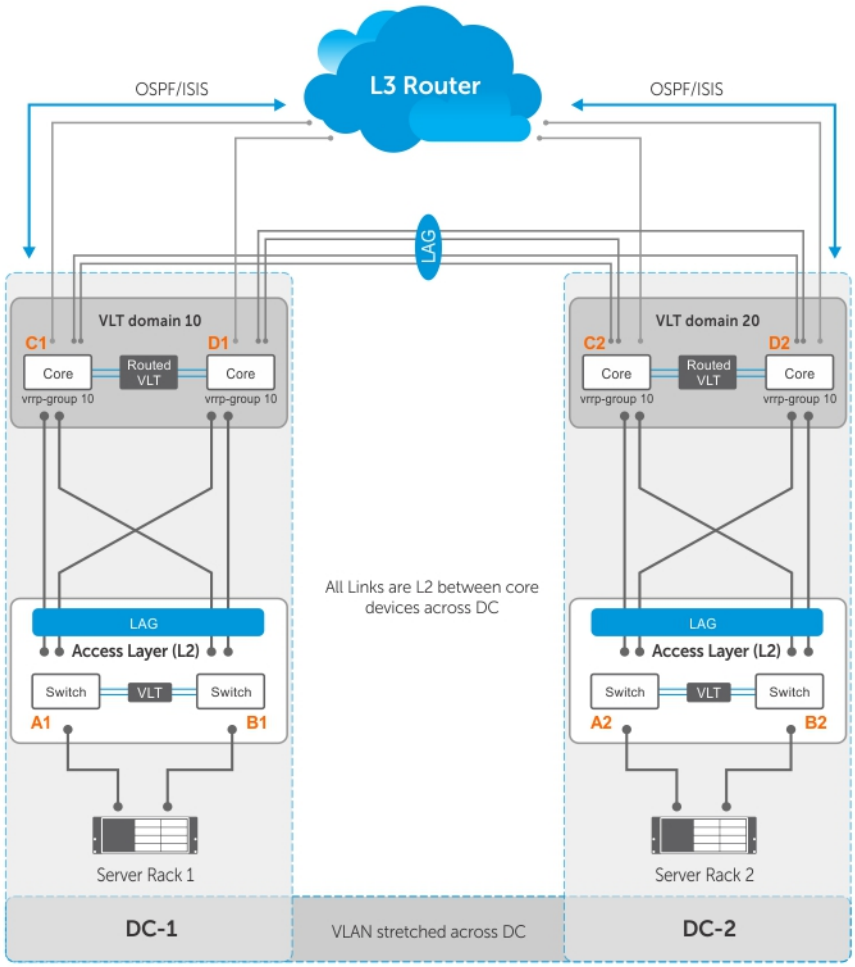
Migrate VMs across data centers

OS10 switches support movement of virtual machines (VMs) across data centers using VRRP Active-Active mode.

Configure symmetric VRRP with same VRRP group ID and virtual IP in VLANs stretched or spanned across data centers. VMs use the VRRP Virtual IP address of the VLAN as Gateway IP. As the VLAN configurations are symmetric across data centers, you can move the VMs from one data center to another..

You must assign the same VRRP group IDs to the VLANs in L3 mode, with VRRP in Active-Active mode.

The following illustration shows a sample configuration with two data centers:



- Server racks, Rack 1 and Rack 2, are part of data centers DC1 and DC2, respectively.
- Rack 1 is connected to devices A1 and B1 in L2 network segment.
- Rack 2 is connected to devices A2 and B2 in L2 network segment.
- A VLT LAG is present between A1 and B1 as well as A2 and B2.
- A1 and B1 connect to core routers, C1 and D1 with VLT routing enabled.
- A2 and B2 connect to core routers, C2 and D2, with VLT routing enabled.
- The data centers are connected through a direct link or eVLT.
- The core routers C1 and D1 in the local VLT domain connect to the core routers C2 and D2 in the remote VLT domain using VLT links.
- The core routers C1 and D1 in local VLT domain along with C2 and D2 in the remote VLT domain are part of an L3 cloud.
- The core routers C1, D1, C2, D2 are in a VRRP group with the same vrrp-group ID.

When a virtual machine running in Server Rack 1 migrates to Server Rack 2, L3 packets for that VM are routed without interruption.

Sample configuration of C1:

- **Configure VRRP on L2 links between core routers:**

```
C1(config)# interface vlan 100
C1(conf-if-vl-100)# ip address 10.10.100.1/24
C1(conf-if-vl-100)# vrrp-group 10
C1(conf-vlan100-vrid-10)# priority 250
C1(conf-vlan100-vrid-10)# virtual-address 10.10.100.5
```

- **Configure VLT port channel for VLAN 100:**

```
C1(config)# interface port-channel 10
C1(conf-if-po-10)# vlt-port-channel 10
C1(conf-if-po-10)# switchport mode trunk
C1(conf-if-po-10)# switchport trunk allowed vlan 100
C1(conf-if-po-10)# exit
```

- **Add members to port channel 10:**

```
C1(config)# interface ethernet 1/1/3
C1(conf-if-eth1/1/3)# channel-group 10
C1(conf-if-eth1/1/3)# exit
C1(config)# interface ethernet 1/1/4
C1(conf-if-eth1/1/4)# channel-group 10
C1(conf-if-eth1/1/4)# exit
```

- **Configure OSPF on L3 side of core router:**

```
C1(config)# router ospf 100
C1(conf-router-ospf-100)# exit
C1(config)# interface vlan 200
C1(conf-if-vl-200)# ip ospf 100 area 0.0.0.0
```

- **Configure VLT port channel for VLAN 200:**

```
C1(config)# interface port-channel 20
C1(conf-if-po-20)# vlt-port-channel 20
C1(conf-if-po-20)# switchport mode trunk
C1(conf-if-po-20)# switchport trunk allowed vlan 200
C1(conf-if-po-20)# exit
```

- **Add members to port channel 20:**

```
C1(config)# interface ethernet 1/1/5
C1(conf-if-eth1/1/5)# channel-group 20
C1(conf-if-eth1/1/5)# exit
C1(config)# interface ethernet 1/1/6
C1(conf-if-eth1/1/6)# channel-group 20
C1(conf-if-eth1/1/6)# exit
```

Sample configuration of D1:

- **Configure VRRP on L2 links between core routers:**

```
D1(config)# interface vlan 100
D1(conf-if-vl-100)# ip address 10.10.100.2/24
D1(conf-if-vl-100)# vrrp-group 10
D1(conf-vlan100-vrid-10)# virtual-address 10.10.100.5
```

- **Configure VLT port channel for VLAN 100:**

```
D1(config)# interface port-channel 10
D1(conf-if-po-10)# vlt-port-channel 10
D1(conf-if-po-10)# switchport mode trunk
D1(conf-if-po-10)# switchport trunk allowed vlan 100
D1(conf-if-po-10)# exit
```

- **Add members to port channel 10:**

```
D1(config)# interface ethernet 1/1/3
D1(conf-if-eth1/1/3)# channel-group 10
D1(conf-if-eth1/1/3)# exit
D1(config)# interface ethernet 1/1/4
D1(conf-if-eth1/1/4)# channel-group 10
D1(conf-if-eth1/1/4)# exit
```

- **Configure OSPF on L3 side of core router:**

```
D1(config)# router ospf 100
D1(conf-router-ospf-100)# exit
D1(config)# interface vlan 200
D1(conf-if-vl-200)# ip ospf 100 area 0.0.0.0
```

- **Configure VLT port channel for VLAN 200:**

```
D1(config)# interface port-channel 20
D1(conf-if-po-20)# vlt-port-channel 20
D1(conf-if-po-20)# switchport mode trunk
```

```
D1(conf-if-po-20)# switchport trunk allowed vlan 200
D1(conf-if-po-20)# exit
```

- **Add members to port channel 20:**

```
D1(config)# interface ethernet 1/1/5
D1(conf-if-eth1/1/5)# channel-group 20
D1(conf-if-eth1/1/5)# exit
D1(config)# interface ethernet 1/1/6
D1(conf-if-eth1/1/6)# channel-group 20
D1(conf-if-eth1/1/6)# exit
```

Sample configuration of C2:

- **Configure VRRP on L2 links between core routers:**

```
C2(config)# interface vlan 100
C2(conf-if-vl-100)# ip address 10.10.100.3/24
C2(conf-if-vl-100)# vrrp-group 10
C2(conf-vlan100-vrid-10)# virtual-address 10.10.100.5
```

- **Configure VLT port channel for VLAN 100:**

```
C2(config)# interface port-channel 10
C2(conf-if-po-10)# vlt-port-channel 10
C2(conf-if-po-10)# switchport mode trunk
C2(conf-if-po-10)# switchport trunk allowed vlan 100
C2(conf-if-po-10)# exit
```

- **Add members to port channel 10:**

```
C2(config)# interface ethernet 1/1/3
C2(conf-if-eth1/1/3)# channel-group 10
C2(conf-if-eth1/1/3)# exit
C2(config)# interface ethernet 1/1/4
C2(conf-if-eth1/1/4)# channel-group 10
C2(conf-if-eth1/1/4)# exit
```

- **Configure OSPF on L3 side of core router:**

```
C2(config)# router ospf 100
C2(conf-router-ospf-100)# exit
C2(config)# interface vlan 200
C2(conf-if-vl-200)# ip ospf 100 area 0.0.0.0
```

- **Configure VLT port channel for VLAN 200:**

```
C2(config)# interface port-channel 20
C2(conf-if-po-20)# vlt-port-channel 20
C2(conf-if-po-20)# switchport mode trunk
C2(conf-if-po-20)# switchport trunk allowed vlan 200
C2(conf-if-po-20)# exit
```

- **Add members to port channel 20:**

```
C2(config)# interface ethernet 1/1/5
C2(conf-if-eth1/1/5)# channel-group 20
C2(conf-if-eth1/1/5)# exit
C2(config)# interface ethernet 1/1/6
C2(conf-if-eth1/1/6)# channel-group 20
C2(conf-if-eth1/1/6)# exit
```

Sample configuration of D2:

- **Configure VRRP on L2 links between core routers:**

```
D2(config)# interface vlan 100
D2(conf-if-vl-100)# ip address 10.10.100.4/24
D2(conf-if-vl-100)# vrrp-group 10
D2(conf-vlan100-vrid-10)# virtual-address 10.10.100.5
```

- **Configure VLT port channel for VLAN 100:**

```
D2(config)# interface port-channel 10
D2(conf-if-po-10)# vlt-port-channel 10
D2(conf-if-po-10)# switchport mode trunk
D2(conf-if-po-10)# switchport trunk allowed vlan 100
D2(conf-if-po-10)# exit
```

- **Add members to port channel 10:**

```
D2(config)# interface ethernet 1/1/3
D2(conf-if-eth1/1/3)# channel-group 10
D2(conf-if-eth1/1/3)# exit
D2(config)# interface ethernet 1/1/4
D2(conf-if-eth1/1/4)# channel-group 10
D2(conf-if-eth1/1/4)# exit
```

- **Configure OSPF on L3 side of core router:**

```
D2(config)# router ospf 100
D2(conf-router-ospf-100)# exit
D2(config)# interface vlan 200
D2(conf-if-vl-200)# ip ospf 100 area 0.0.0.0
```

- **Configure VLT port channel for VLAN 200:**

```
D2(config)# interface port-channel 20
D2(conf-if-po-20)# vlt-port-channel 20
D2(conf-if-po-20)# switchport mode trunk
D2(conf-if-po-20)# switchport trunk allowed vlan 200
D2(conf-if-po-20)# exit
```

- **Add members to port channel 20:**

```
D2(config)# interface ethernet 1/1/5
D2(conf-if-eth1/1/5)# channel-group 20
D2(conf-if-eth1/1/5)# exit
D2(config)# interface ethernet 1/1/6
D2(conf-if-eth1/1/6)# channel-group 20
D2(conf-if-eth1/1/6)# exit
```

View VLT information

To monitor the operation or verify the configuration of a VLT domain, use a VLT `show` command on primary and secondary peers.

- View detailed information about the VLT domain configuration in EXEC mode, including VLTi status, local and peer MAC addresses, peer-routing status, and VLT peer parameters.

```
show vlt domain-id
```

- View the role of the local and remote VLT peer in EXEC mode.

```
show vlt domain-id role
```

- View any mismatches in the VLT configuration in EXEC mode.

```
show vlt domain-id mismatch
```

- View detailed information about VLT ports in EXEC mode.

```
show vlt domain-id vlt-port-detail
```

- View the current configuration of all VLT domains in EXEC mode.

```
show running-configuration vlt
```

View peer-routing information

```
OS10# show vlt 255
Domain ID                  : 255
Unit ID                    : 1
Role                       : primary
Version                    : 2.0
Local System MAC address   : 34:17:eb:3a:bd:80
Role priority              : 1
VLT MAC address            : aa:bb:cc:dd:ee:ff
IP address                 : fda5:74c8:b79e:1::1
Delay-Restore timer        : 100 seconds
Peer-Routing               : Enabled
Peer-Routing-Timeout timer : 9999 seconds
VLTi Link Status
    port-channel1000       : up

VLT Peer Unit ID    System MAC Address    Status    IP Address              Version
--------------------------------------------------------------------------------
```

```
  2                 34:17:eb:3a:c2:80      up       fda5:74c8:b79e:1::2     2.0
OS10#
```

View VLT role

* indicates the local peer

```
OS10# show vlt 1 role
VLT Unit ID    Role
------------------------
* 1            primary
  2            secondary
```

View VLT mismatch — no mismatch

```
OS10# show vlt 1 mismatch
Peer-routing mismatch:
No mismatch

VLAN mismatch:
No mismatch

VLT VLAN mismatch:
No mismatch
```

View VLT mismatch — mismatch in VLT configuration

```
OS10# show vlt 1 mismatch peer-routing
Peer-routing mismatch:
VLT Unit ID         Peer-routing
-----------------------------
* 1                  Enabled
  2                  Disabled


OS10# show vlt 1 mismatch
Peer-routing mismatch:
VLT Unit ID        Peer-routing
----------------------------
* 1                Enabled
  2                Disabled

VLAN mismatch:
VLT Unit ID        Mismatch VLAN List
----------------------------------
* 1                -
  2                4

VLT VLAN mismatch:
VLT ID : 1
VLT Unit ID        Mismatch VLAN List
--------------------------------
* 1              1
  2              2
VLT ID : 2
VLT Unit ID    Mismatch VLAN List
----------------------------------
* 1               1
  2               2
```

View VLT port details

* indicates the local peer

```
OS10# show vlt 1 vlt-port-detail
VLT port channel ID : 1
VLT Unit ID    Port-Channel    Status   Configured ports  Active ports
-------------------------------------------------------------------------
```

```
* 1            port-channel1   down     2                 0
  2            port-channel1   down     2                 0
VLT port channel ID : 2
VLT Unit ID    Port-Channel    Status   Configured ports  Active ports
--------------------------------------------------------------------
* 1            port-channel2   down     1                 0
  2            port-channel2   down     1                 0
VLT port channel ID : 3
VLT Unit ID    Port-Channel    Status   Configured ports  Active ports
--------------------------------------------------------------------
  2            port-channel3   down     1                 0
```

View VLT running configuration

```
OS10# show running-configuration vlt
!
vlt domain 1
 peer-routing
 discovery-interface ethernet1/1/17
!
interface port-channel1
 vlt-port-channel 10
!
interface port-channel10
 vlt-port-channel 20
!
interface port-channel20
 vlt-port-channel 20
```

VLT commands

backup destination

Configures the VLT backup link for heartbeat timers.

Syntax `backup destination {ip-address | ipv6 ipv6-address} [vrf management] [interval interval-time]`

Parameters

- *ip-address* — Enter the IPv4 address of the backup link.
- *ipv6-address* — Enter the IPv6 address of the backup link.
- `vrf management` — (Optional) Configure the management VRF instance for the backup IPv4 or IPv6 address.
- `interval` *interval-time* — (Optional) Enter the time in seconds to configure the heartbeat interval.

Default Not configured

Command Mode VLT-DOMAIN

Usage Information The `no` version of this command removes the IP address from the backup link.

Example

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination 10.16.151.110 vrf management interval 30
```

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# backup destination ipv6 1::1 vrf management interval 30
```

Supported Releases 10.3.1E or later

delay-restore

Configures a time interval to delay VLT ports bring up after reload or peer-link restoration between the VLT peer switches.

Syntax `delay-restore seconds`

Parameters *seconds* — Enter a delay time, in seconds, to delay bringing up VLT ports after the VLTi device is reloaded, from 1 to 1200.

Default 90 seconds

Command Mode VLT-DOMAIN

Usage Information Use this command to delay the system from bringing up the VLT port for a brief period to allow L3 routing protocols to converge. If the peer VLT device was up at the time the VLTi link failed, use this command after you reload a VLT device. The `no` version of this command resets the delay time to the default value.

Example

```
OS10(conf-vlt-1)# delay-restore 100
```

Supported Releases 10.3.0E or later

discovery-interface

Configures the interface to discover and connect to a VLT peer in the VLT interconnect (VLTi) link between peers.

Syntax `discovery-interface {ethernet node/slot/port[:subport]}`

Parameters *ethernet* — Enter the Ethernet interface information for the port on a VLT peer. You can also enter a range of interfaces separated by hyphens.

Default None

Command Mode VLT-DOMAIN

Usage Information The VLT node discovery service auto-LAGs the discovery ports and creates VLTi interfaces. The `no` version of this command disables the discovery-interface configuration.

Example

```
OS10(config)# vlt-domain 1
OS10(conf-vlt-1)# discovery-interface ethernet 1/1/15
```

Example (range)

```
OS10(config)# vlt-domain 2
OS10(conf-vlt-2)# discovery-interface ethernet 1/1/1-1/1/12
```

Supported Releases 10.2.0E or later

peer-routing

Enables or disables L3 routing to peers.

Syntax `peer-routing`

Parameters None

Default Disabled

Command Mode VLT-DOMAIN

Usage Information	The `no` version of this command disables L3 routing.
Example	`OS10(conf-vlt-1)# peer-routing`
Supported Releases	10.2.0E or later

peer-routing-timeout

Configures the delay after which peer routing disables when the peer is not available. This command supports both IPv6 and IPv4.

Syntax	`peer-routing-timeout` *value*
Parameters	*value* — Enter the timeout value in seconds, from 0 to 65535.
Default	0
Command Mode	VLT-DOMAIN
Usage Information	When the timer expires, the software checks to see if the VLT peer is available. If the VLT peer is not available, peer-routing disables on the peer. If you do not configure the timer, peer-routing does not disable even when the peer is unavailable.
Example	`OS10(conf-vlt-1)# peer-routing-timeout 120`
Supported Releases	10.3.0E or later

primary-priority

Configures the priority when selecting the primary and secondary VLT peers during election.

Syntax	`primary-priority` *value*
Parameters	*value* — Enter a lower value than the priority value of the remote peer. The range is from 1 to 65535. The default value is 32768.
Default	32768.
Command Mode	VLT-DOMAIN
Usage Information	• After you configure a VLT domain on each peer switch and connect the two VLT peers on each side of the VLT interconnect, the system elects a primary and secondary VLT peer device. To configure the primary and secondary roles before the election process, use the `primary-priority` command. Enter a lower value on the primary peer and a higher value on the secondary peer. If the primary peer fails, the secondary peer (with the higher priority) takes the primary role. If the primary peer (with the lower priority) later comes back online, it is assigned the secondary role; there is no preemption. • If the priority values configured on the two VLT peers are equal, VLT uses the default primary election mechanism based on the values of the system MAC addresses of the two nodes. The VLT peer with the lowest system MAC address assumes the primary role. • If the heartbeat is up and the VLTi link goes down between the VLT peers, both the VLT peers retain their primary and secondary roles. However, the VLT LAG on the secondary VLT peer shuts down. **NOTE: When you configure a priority for VLT peers using this command, the configuration is not effective immediately. The primary priority configuration comes into effect the next time election is triggered.**

Example

```
OS10(conf-vlt-1)#primary-priority 2
```

Supported Releases 10.4.1.0 or later

show spanning-tree virtual-interface

Displays STP , RPVST+, and MSTP information specific to VLT.

Syntax `show spanning-tree virtual-interface [detail]`

Parameters `detail`—(Optional) Displays detailed output.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show spanning-tree virtual-interface
VFP(VirtualFabricPort) of RSTP 1 is Designated Forwarding
Edge port: No (default)
Link type: point-to-point (auto)
Boundary: No, Bpdu-filter: Disable, Bpdu-Guard: Disable, Shutdown-on-Bpdu-Guard-violation: No
Root-Guard: Disable, Loop-Guard: Disable
Bpdus (MRecords) Sent: 11, Received: 7
Interface                                                                    Designated
Name                    PortID    Prio   Cost   Sts    Cost  Bridge ID      PortID
-------------------------------------------------------------------------------------------
VFP(VirtualFabricPort)  0.1       0      1      FWD    0     32768    0078.7614.6062  0.1
```

```
OS10# show spanning-tree virtual-interface

VFP(VirtualFabricPort) of vlan 100 is Designated Blocking
Edge port: No (default)
Link type: point-to-point (auto)
Boundary: No, Bpdu-filter: Disable, Bpdu-Guard: Disable, Shutdown-on-Bpdu-Guard-violation: No
Root-Guard: Disable, Loop-Guard: Disable
Bpdus (MRecords) Sent: 7, Received: 9
Interface                                                          Designated
Name                     PortID   Prio    Cost   Sts  Cost  Bridge ID   PortID
-------------------------------------------------------------------------------------------
VFP(VirtualFabricPort)   0.1       0      1      BLK   0    4196      90b1.1cf4.a602  0.1
```

Example (detail)

```
OS10# show spanning-tree virtual-interface detail
Port 1 (VFP(VirtualFabricPort)) of RSTP 1 is designated Forwarding
Port path cost 1, Port priority 0, Port Identifier 0.1
Designated root priority: 32768, address: 00:78:76:14:60:62
Designated bridge priority: 32768, address: 00:78:76:14:60:62
Designated port ID: 0.1, designated path cost: 0
Number of transitions to forwarding state: 1
Edge port: No (default)
Link Type: Point-to-Point
BPDU Sent: 15, Received: 5
```

```
OS10# show spanning-tree virtual-interface detail
Port 1 (VFP(VirtualFabricPort)) of vlan1 is designated Forwarding
Port path cost 1, Port priority 0, Port Identifier 0.1
Designated root priority: 4097, address: 90:b1:1c:f4:a6:02
Designated bridge priority: 4097, address: 90:b1:1c:f4:a6:02
Designated port ID: 0.1, designated path cost: 0
Number of transitions to forwarding state: 1
Edge port: No (default)
Link Type: Point-to-Point
BPDU Sent: 202, Received: 42
Port 1 (VFP(VirtualFabricPort)) of vlan100 is designated Forwarding
Port path cost 1, Port priority 0, Port Identifier 0.1
Designated root priority: 4196, address: 90:b1:1c:f4:a6:02
Designated bridge priority: 4196, address: 90:b1:1c:f4:a6:02
Designated port ID: 0.1, designated path cost: 0
Number of transitions to forwarding state: 1
```

```
Edge port: No (default)
Link Type: Point-to-Point
BPDU Sent: 101, Received: 21
```

Example (MSTP information)

```
OS10# show spanning-tree virtual-interface
VFP(VirtualFabricPort) of MSTI 0 is Designated Forwarding
Edge port: No (default)
Link type: point-to-point (auto)
Boundary: Yes, Bpdu-filter: Disable, Bpdu-Guard: Disable, Shutdown-on-Bpdu-Guard-violation: No
Root-Guard: Disable, Loop-Guard: Disable
Bpdus (MRecords) Sent: 387, Received: 16
Interface                                                                    Designated
Name                PortID    Prio      Cost      Sts         Cost      Bridge ID                PortID
-------------------------------------------------------------------------------------------------------------
VFP(VirtualFabricPort) 0.1     0         1         FWD         0         32768    3417.ebf2.a8c4  0.1

VLT-LAG -1(vlt-portid-1) of MSTI 0 is in Designated Forwarding
Edge port: No (default)
Link type: point-to-point (auto)
Boundary: No, Bpdu-filter: Disable, Bpdu-Guard: Disable,
Shutdown-on-Bpdu-Guard-violation: No
Root-Guard: Disable, Loop-Guard: Disable
Bpdus (MRecords) Sent: 1234, Received: 123

Virtual
Interface                                                          Designated
Name              PortID    Prio      Cost      Sts         Cost    Bridge ID                PortID
-------------------------------------------------------------------------------------------------------

VLT-LAG -1(vlt-portid1)         128.2001  128       2000000   FWD         0       32768    90b1.1cf4.a523  128.2001
```

Example (MSTP information on VLT)

```
OS10# show spanning-tree virtual-interface detail
Port 1 (VFP(VirtualFabricPort)) of MSTI 0 is designated Forwarding
Port path cost 0, Port priority 128, Port Identifier 128.1
Designated root priority: 32768, address: 34:17:44:55:66:7f
Designated bridge priority: 32768, address: 90:b1:1c:f4:a5:23
Designated port ID: 128.1, designated path cost: 0
Number of transitions to forwarding state: 1
Edge port: No (default)
Link Type: Point-to-Point
BPDU Sent: 2714, Received: 1234

Port 2001 (VLT-LAG -1(vlt-portid-1)) of MSTI 0 is designated Forwarding
Port path cost 200000, Port priority 128, Port Identifier 128.2001
Designated root priority: 32768, address: 34:17:44:55:66:7f
Designated bridge priority: 32768, address: 90:b1:1c:f4:a5:23
Designated port ID: 128.2001, designated path cost: 0
Number of transitions to forwarding state: 1
Edge port: No (default)
Link Type: Point-to-Point
BPDU Sent: 2714, Received: 1234
```

Supported Releases 10.3.0E or later

show vlt

Displays information on a VLT domain.

Syntax `show vlt id`

Parameter *id* — Enter a VLT domain ID, from 1 to 255.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show vlt 255
Domain ID                  : 255
Unit ID                    : 1
Role                       : primary
```

```
Version                     : 2.0
Local System MAC address    : 34:17:eb:3a:bd:80
Role priority               : 1
VLT MAC address             : aa:bb:cc:dd:ee:ff
IP address                  : fda5:74c8:b79e:1::1
Delay-Restore timer         : 100 seconds
Peer-Routing                : Enabled
Peer-Routing-Timeout timer : 9999 seconds
VLTi Link Status
    port-channel1000        : up

VLT Peer Unit ID    System MAC Address    Status    IP Address              Version
--------------------------------------------------------------------------------
  2                 34:17:eb:3a:c2:80      up       fda5:74c8:b79e:1::2     2.0
```

Supported Releases 10.2.0E or later

show vlt backup-link

Displays the details of heartbeat status.

Syntax `show vlt domain-id backup-link`

Parameters *domain-id* — Enter the VLT domain ID.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show vlt 255 backup-link
VLT Backup Link
------------------------
Destination                    : 10.16.208.164
Peer Heartbeat status          : Up
Heartbeat interval             : 1
Heartbeat timeout              : 3
```

Supported Releases 10.3.1E or later

show vlt mac-inconsistency

Displays inconsistencies in dynamic MAC addresses learnt between VLT peers across spanned-vlans.

Syntax `show vlt mac-inconsistency`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information Use this command to check for a mismatch of MAC address table entries between VLT peers. To verify VLT configuration mismatch issues on peer switches, use the `show vlt domain-name mismatch command`.

Example

```
OS10# show vlt-mac-inconsistency
Checking Vlan 228 .. Found 7 inconsistencies .. Progress 100%
VLAN 128
----------
MAC 00:00:00:00:00:02 is missing from Node(s) 2
VLAN 1
```

```
----------
MAC 00:a0:c9:00:00:18 is missing from Node(s) 2
MAC 00:a0:c9:00:00:20 is missing from Node(s) 2
VLAN 131
----------
MAC 00:00:00:00:00:02 is missing from Node(s) 2
VLAN 132
----------
MAC 00:00:00:00:00:02 is missing from Node(s) 2
VLAN 135
----------
MAC 00:00:00:00:00:02 is missing from Node(s) 2
VLAN 137
----------
MAC 00:00:00:00:00:02 is missing from Node(s) 2

Run "show vlt d1 mismatch ..." commands to identify configuration issues
```

Supported Releases 10.2.0E or later

show vlt mismatch

Displays mismatches in a VLT domain configuration.

Syntax `show vlt id mismatch [peer-routing | vlan | vlt-vlan vlt-port-id | virtual-network]`

Parameters

- *id* — Enter the VLT domain ID, from 1 to 255.
- `peer-routing` — Display mismatches in peer-routing configuration.
- `vlan` — Display mismatches in a VLAN configuration in the VLT domain.
- `vlt-vlan vlt-port-id` — Display mismatches in VLT port configuration, from 1 to 4095.
- `virtual-network` — Display mismatches in virtual network configurations between VLT peers.

Default Not configured

Command Mode EXEC

Usage Information The * in the mismatch output indicates a local node entry.

Example (no mismatch)

```
OS10# show vlt 1 mismatch
Peer-routing mismatch:
No mismatch

VLAN mismatch:
No mismatch

VLT VLAN mismatch:
No mismatch
```

Example (mismatch)

```
OS10# show vlt 1 mismatch
Peer-routing mismatch:
VLT Unit ID    Peer-routing
-----------------------------------
* 1            Enabled
  2            Disabled

VLAN mismatch:
No mismatch

VLT VLAN mismatch:
VLT ID : 1
```

```
VLT Unit ID    Mismatch VLAN List
---------------------------------
* 1              1
  2              2
VLT ID : 2
VLT Unit ID    Mismatch VLAN List
---------------------------------
* 1              1
  2              2
```

Example (mismatch peer routing)

```
OS10# show vlt 1 mismatch peer-routing
Peer-routing mismatch:
VLT Unit ID    Peer-routing
-----------------------------------
* 1            Enabled
  2            Disabled
```

Example (mismatch VLAN)

```
OS10# show vlt 1 mismatch vlan
VLT Unit ID    Mismatch VLAN List
---------------------------------
* 1              -
  2              4
```

Example (mismatch VLT VLAN)

```
OS10# show vlt 1 mismatch vlt-vlan
VLT ID : 1
VLT Unit ID    Mismatch VLAN List
-------------------------------------
* 1              1
  2              2
VLT ID : 2
VLT Unit ID    Mismatch VLAN List
---------------------------------
* 1              1
  2              2
```

Example (mismatch — Virtual Network (VN) name not available in the peer)

```
OS10# show vlt all mismatch virtual-network
Virtual Network Name Mismatch:
VLT Unit ID    Mismatch Virtual Network List
-------------------------------------------------------------------------------
  1              10,104
* 2              -
```

Example (mismatch of VLTi and VLAN)

```
OS10# show vlt all mismatch virtual-network
Virtual Network: 100
VLT Unit ID    Configured VLTi-Vlans
-------------------------------------------------------------------------------
  1              101
* 2              100
```

Example (mismatch of VN mode)

```
OS10# show vlt all mismatch virtual-network
Virtual Network: 102
VLT Unit ID    Configured Virtual Network Mode
-------------------------------------------------------------------------------
  1              PV
* 2              Attached
```

Example (mismatch of port and VLAN list)

```
OS10# show vlt all mismatch virtual-network
Virtual Network: 102
VLT Unit ID    Mismatch (VLT Port,Vlan) List
-------------------------------------------
  1              -
* 2              (vlt-port-channel10,vlan99)
```

```
Virtual Network: 103
VLT Unit ID    Mismatch (VLT Port,Vlan) List
-----------------------------------------
  1              (vlt-port-channel10,vlan103)
* 2              (vlt-port-channel10,vlan104)
```

Example (mismatch of untagged interfaces)

```
OS10# show vlt all mismatch virtual-network
Virtual Network: 104
VLT Unit ID    Mismatch Untagged VLT Port-channel List
-------------------------------------------------------
  1              10
* 2              -
```

Supported Releases 10.2.0E or later

show vlt role

Displays the VLT role of the local peer.

Syntax `show vlt id role`

Parameters `id` — Enter the VLT domain ID, from 1 to 255.

Default Not configured

Command Mode EXEC

Usage Information The * in the mismatch output indicates a local node entry.

Example

```
OS10# show vlt 1 role
VLT Unit ID    Role
------------------------
* 1            primary
  2            secondary
```

Supported Releases 10.2.0E or later

show vlt vlt-port-detail

Displays detailed status information about VLT ports.

Syntax `show vlt id vlt-port-detail`

Parameters *id* — Enter a VLT domain ID, from 1 to 255.

Default Not configured

Command Mode EXEC

Usage Information The * in the mismatch output indicates a local node entry.

Example

```
OS10# show vlt 1 vlt-port-detail
Vlt-port-channel ID : 1
VLT Unit ID   Port-Channel    Status   Configured ports  Active ports
-------------------------------------------------------------------
* 1           port-channel1   down     2                 0
  2           port-channel1   down     2                 0
VLT ID : 2
VLT Unit ID   Port-Channel    Status   Configured ports  Active ports
-------------------------------------------------------------------
* 1           port-channel2   down     1                 0
```

```
  2           port-channel2   down     1                0
VLT ID : 3
VLT Unit ID   Port-Channel    Status   Configured ports  Active ports
------------------------------------------------------------------
  2           port-channel3   down     1                0
```

Supported Releases 10.2.0E or later

vlt-domain

Creates a VLT domain.

Syntax `vlt-domain domain-id`

Parameter *domain-id* — Enter a VLT domain ID on each peer, from 1 to 255.

Default None

Command Mode CONFIGURATION

Usage Information Configure the same VLT domain ID on each peer. If a VLT domain ID mismatch occurs on VLT peers, the VLTi link between peers does not activate. The no version of this command disables VLT.

Example

```
OS10(config)# vlt-domain 1
```

Supported Releases 10.2.0E or later

vlt-port-channel

Configures the ID used to map interfaces on VLT peers into a single VLT port-channel.

Syntax `vlt-port-channel vlt-lag-id`

Parameters *vlt-lag-id* — Enter a VLT port-channel ID, from 1 to 1024.

Default Not configured

Command Mode PORT-CHANNEL INTERFACE

Usage Information Assign the same VLT port-channel ID to interfaces on VLT peers to create a VLT port-channel. The `no` version of this command removes the VLT port-channel ID configuration.

Example (peer 1)

```
OS10(conf-if-po-10)# vlt-port-channel 1
```

Example (peer 2)

```
OS10(conf-if-po-20)# vlt-port-channel 1
```

Supported Releases 10.2.0E or later

vlt-mac

Configures a MAC address for all peer switches in a VLT domain.

Syntax `vlt-mac mac-address`

Parameters *mac-address* — Enter a MAC address for the topology in nn:nn:nn:nn:nn:nn format.

Default Not configured

Command Mode	VLT-DOMAIN
Usage Information	Use this command to minimize the time required to synchronize the default MAC address of the VLT domain on both peer devices when one peer switch reboots. If you do not configure a VLT MAC address, the MAC address of the primary peer is used as the VLT MAC address across all peers. This configuration must be symmetrical in all the peer switches to avoid any unpredictable behavior. For example, unit down or VLTi reset. The `no` version of this command disables the VLT MAC address configuration.

NOTE: Configure the VLT MAC address as symmetrical in all the VLT peer switches to avoid any unpredictable behavior when any unit is down or when VLTi is reset.

Example	`OS10(conf-vlt-1)# vlt-mac 00:00:00:00:00:02`
Supported Releases	10.2.0E or later

vrrp mode active-active

Enables the VRRP peers to locally forward L3 traffic in a VLAN interface.

Syntax	`vrrp mode active-active`
Parameters	None
Default	Enabled
Command Mode	VLAN INTERFACE
Usage Information	This command is applicable only for VLAN interfaces. In a non-VLT network, the backup VRRP gateway forwards L3 traffic. If you want to use VRRP groups on VLANs without VLT topology, disable the Active-Active functionality, to ensure that only the active VRRP gateway forwards L3 traffic. The `no` version of this command disables the configuration.
Example	`OS10(conf-if-vl-10)# vrrp mode active-active`
Supported Releases	10.2.0E or later

14

Uplink Failure Detection

Uplink failure detection (UFD) indicates the loss of upstream connectivity to servers connected to the switch.

A switch provides upstream connectivity for devices, such as servers. If the switch loses upstream connectivity, the downstream devices also lose connectivity. However, the downstream devices do not generally receive an indication that the upstream connectivity was lost because connectivity to the switch is still operational. To solve this issue, use UFD.

UFD associates downstream interfaces with upstream interfaces. When upstream connectivity fails, the switch operationally disables its downstream links. Failures on the downstream links allow downstream devices to recognize the loss of upstream connectivity. This allows the downstream servers to select alternate paths, if available, to send traffic to upstream devices.

UFD creates an association between upstream and downstream interfaces known as *uplink-state group*. An interface in an uplink-state group can be a physical Ethernet or fibre channel interface or a port-channel.

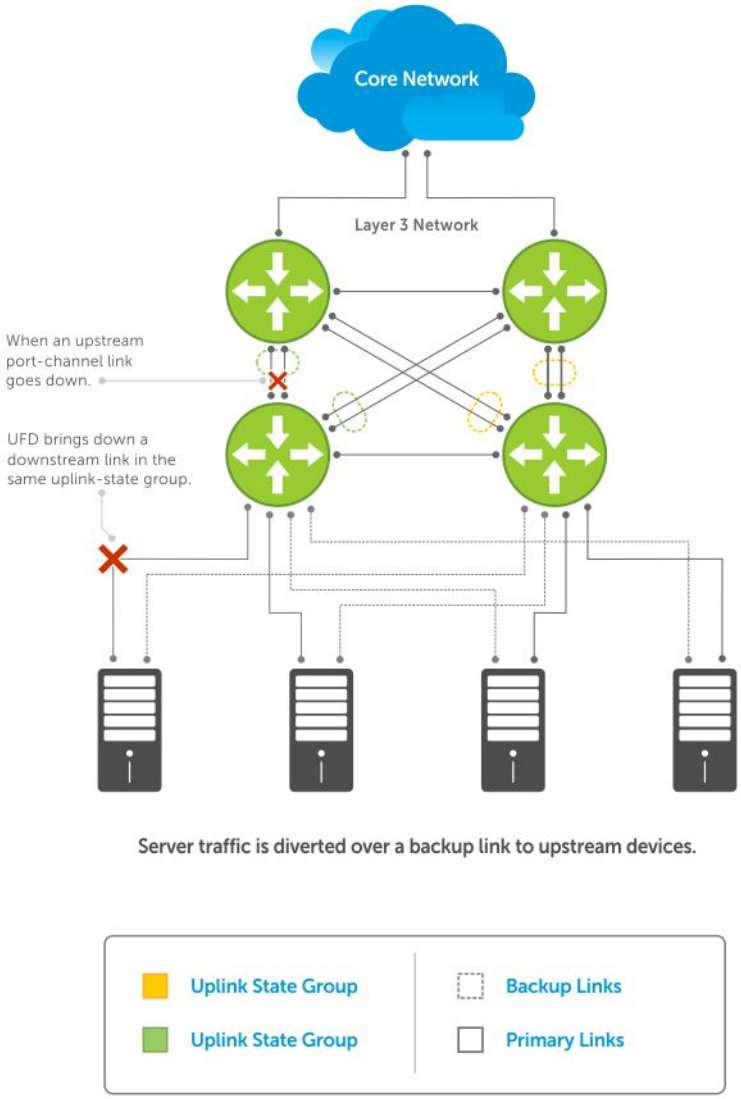
An enabled uplink-state group tracks the state of all assigned upstream interfaces. The failure of upstream interfaces results in automatic disabling of downstream interfaces in the uplink-state group, as shown in the following illustration. If only one of the upstream interfaces in an uplink-state group goes down, a specific number of downstream interfaces in the same uplink-state group go down. You can configure the number of downstream interfaces that go down based on the traffic conditions from the server to the upstream interfaces. This avoids overloading traffic on upstream ports.

By default, if all the upstream interfaces in an uplink-state group go down, all the downstream interfaces in the same uplink-state group are set into a link-down state.

In addition, in an uplink-state group, you can configure automatic recovery of downstream ports when there is a change in the link status of uplink interfaces.

You can also bring up downstream interfaces that are in an UFD-disabled error state manually.

UFD Topology



Server traffic is diverted over a backup link to upstream devices.

Configure uplink failure detection

Consider the following before configuring an uplink-state group:

- You can assign a physical port or a port channel to an uplink-state group.
- You can assign an interface to only one uplink-state group at a time.
- You can designate the uplink-state group as either an upstream or downstream interface, but not both.
- You can configure multiple uplink-state groups and operate them concurrently.
- You cannot assign both a port channel and its members to an uplink-state group, which would make the group inactive. The port channels and individual ports that are not part of any port channel can coexist as members of an uplink-state group.
- If one of the upstream interfaces in an uplink-state group goes down, you can set the downstream ports in an operationally down state with an *UFD Disabled error* status. You can configure the system to disable either a user-configurable set of downstream ports or all the downstream ports in the group.
- The downstream ports are disabled in order starting from the lowest numbered port to the highest numbered port.
- When an upstream interface in an uplink-state group that was down comes up, the set of UFD-disabled downstream ports that were down due to that particular upstream interface are brought up, and the *UFD Disabled error* clears in those downstream ports.
- If you disable an uplink-state group, the downstream interfaces are not disabled, regardless of the state of the upstream interfaces.

- If you do not assign upstream interfaces to an uplink-state group, the downstream interfaces are not disabled.

Configuration:

1 Create an uplink-state group in CONFIGURATION mode.

```
uplink-state-group group-id
```

2 Configure the upstream and downstream interfaces in UPLINK-STATE-GROUP mode.

```
upstream {interface-type | interface-range[ track-vlt-status ] | VLTi}
downstream {interface-type | interface-range}
```

3 (Optional) Disable uplink-state group tracking in UPLINK-STATE-GROUP mode.

```
no enable
```

4 (Optional) Provide a descriptive name for the uplink-state group in UPLINK-STATE-GROUP mode.

```
name string
```

5 Configure the number of downstream interfaces to disable, when an upstream interface goes down in UPLINK-STATE-GROUP mode.

```
downstream disable links{number | all}
```

6 (Optional) Enable auto-recovery of downstream interfaces that are disabled in UPLINK-STATE-GROUP mode.

```
downstream auto-recover
```

7 (Optional) Configure the timer to defer the UFD actions on downstream ports in UPLINK-STATE-GROUP mode. When you have configured to track the VLT status in a VLT network, if VLT port-channel is an upstream member of uplink-state group, then the defer timer triggers when the VLT status goes operationally down instead of the operational status of the peer port-channel.

```
defer-time timer
```

8 (Optional) Clear the UFD error disabled state of downstream interfaces in EXEC mode.

```
clear ufd-disable
```

Configure uplink state group

```
OS10(config)# uplink-state-group 1
OS10(conf-uplink-state-group-1)# upstream ethernet 1/1/7:1
OS10(conf-uplink-state-group-1)# downstream ethernet 1/1/1-1/1/5
OS10(conf-uplink-state-group-1)# downstream ethernet 1/1/9:2-1/1/9:3
OS10(conf-uplink-state-group-1)# enable
OS10(conf-uplink-state-group-1)# name UFDGROUP1
OS10(conf-uplink-state-group-1)# defer-time 10
OS10(conf-uplink-state-group-1)# no downstream auto-recover
OS10(conf-uplink-state-group-1)# downstream disable links 2
```

View uplink state group configuration

```
OS10#show uplink-state-group 1

Uplink State Group: 1 Status: Enabled,down
```

```
OS10# show uplink-state-group 1 detail

(Up): Interface up   (Dwn): Interface down   (Dis): Interface disabled

Uplink State Group   :  1     Status   :  Enabled,up Name :  UFDGROUP1
Defer Time           :  10 second(s)
Upstream Interfaces  :  Eth 1/1/7:1(Up)
Downstream Interfaces:  Eth 1/1/1(Dwn)   Eth 1/1/2(Dwn)   Eth 1/1/3(Dwn)   Eth 1/1/4(Dwn)
                        Eth 1/1/5(Dwn)   Eth 1/1/9:2(Dwn) Eth 1/1/9:3(Dwn)
```

```
OS10#show uplink-state-group 1 detail

(Up): Interface up (Dwn): Interface down (Dis): Interface disabled (NA): Not Available

*: VLT port-channel, V: VLT status, P: Peer Operational status ^: Tracking status

Uplink State Group : 1 Name: iscsi_group, Status: Enabled, Up
```

```
Upstream Interfaces : eth1/1/35(Up) *po10(V:Up, ^P:Dwn) VLTi(NA)
Downstream Interfaces : eth1/1/2(Up) *po20(V: Up,P: Up)
```

```
OS10#show uplink-state-group 2 detail

(Up): Interface up (Dwn): Interface down (Dis): Interface disabled (NA): Not Available

*: VLT port-channel, V: VLT status, P: Peer Operational status ^: Tracking status

Uplink State Group : 1 Name: iscsi_group, Status: Enabled, Up
Upstream Interfaces : eth1/1/36(Up) *po30(^V:Up, P:Dwn) VLTi(Up)
Downstream Interfaces : eth1/1/4(Up) *po20(V: Up,P: Up)
```

```
OS10(conf-uplink-state-group-1)# show configuration
!
uplink-state-group 1
 downstream ethernet1/1/1-1/1/5
 downstream ethernet1/1/9:2-1/1/9:3
 upstream ethernet1/1/7:1
```

Uplink failure detection on VLT

When you create uplink-state group in a switch operating in VLT mode, ensure that all the nodes in the VLT setup have same configuration for uplink state groups with VLT port-channel as member. If both the VLT peers do not have the same UFD configuration, the UFD does not work properly.

When you configure VLT port-channel as upstream member in the uplink state group and configure to track the VLT status, the system tracks the fabric Status of VLT. When the fabric status goes down, the uplink state group in each VLT node disables the downstream VLT port-channel local to the node.

When you configure to track the VLT status, the system places the downstream members of the Uplink State Group in error disabled state or clears them from the error disabled state based on the operational status of the VLT port-channel.

When you do not track the VLT status, the system tracks the operational status of port-channel.

Track the VLT status using the `upstream` *`interface-type`* `track-vlt-status` command in UPLINK-STATE-GROUP mode.

To configure VLTi link as member of Uplink State Group, use the `upstream VLTi` command in UPLINK-STATE-GROUP mode. You cannot configure VLTi Link as downstream member in an uplink-state group as UFD may disable the VLTi Link when the upstream members are operationally down. You cannot track the VLT status for an upstream VLTi member.

The following table describes various scenarios when you apply UFD on a VLT network:

Table 69. UFD on VLT network

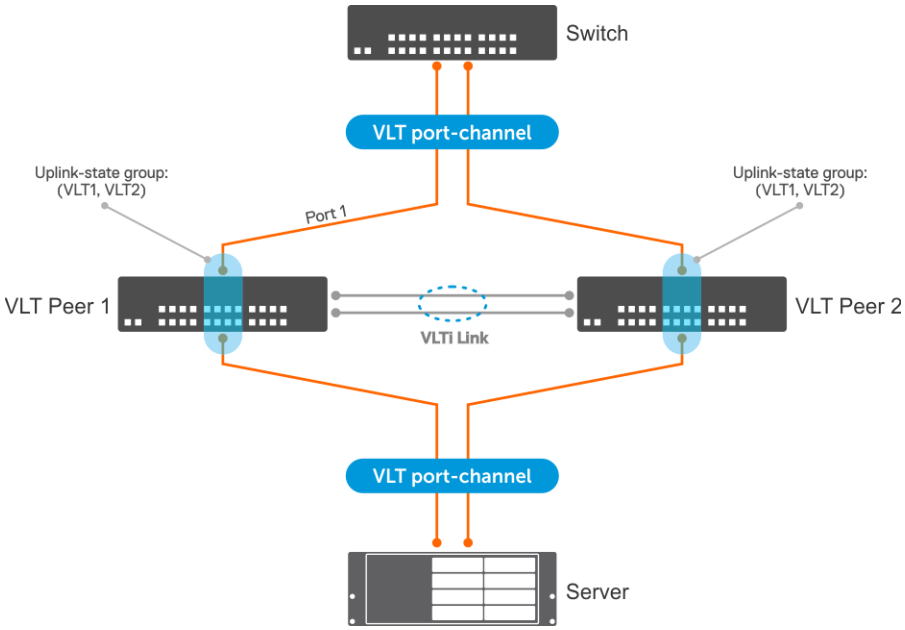
Event	VLT action on primary node	VLT action on secondary node	UFD action
VLTi Link is operationally down with heartbeat up	No action	VLT module sends VLT port-channel disable request to Interface Manager (IFM) for both uplink and downlink.	UFD receives operationally down of upstream VLT port-channel and sends error-disable of downstream VLT port-channel to IFM.
VLTi Link is operationally up with heartbeat up	No action	VLT module sends VLT port-channel enable request to Interface Manager (IFM) for both uplink and downlink.	UFD receives operationally up of upstream VLT port-channel and sends clear error-disable of downstream VLT port-channel to IFM.

Event	VLT action on primary node	VLT action on secondary node	UFD action
Reboot of VLT secondary peer	No action	After reboot, runs the delay restore timer. Both the upstream and downstream VLT port-channel remains disabled until the timer expires.	UFD `error-disables` the downstream VLT port-channel as the upstream VLT port-channel is operationally down. After the timer expires, UFD receives operationally up of upstream VLT port-channel and sends `clear error-disable` of downstream VLT port-channel to IFM.
Reboot of VLT primary peer	Primary becomes secondary peer and runs delay restore timer	Secondary becomes primary	UFD `error-disables` the downstream VLT port-channel as the upstream VLT port-channel is operationally down. After the timer expires, UFD receives operationally up of upstream VLT port-channel and sends `clear error-disable` of downstream VLT port-channel to IFM.
Discovery interface added to UFD group	Invalid configuration	Invalid configuration	Invalid configuration
UFD group member configured as discovery interface	Invalid configuration	Invalid configuration	Invalid configuration
UFD group member made as VLT port-channel	No action	No action	UFD uses fabric status to track the UFD group status.
VLT port-channel added as member of UFD group	No action	No action	UFD uses fabric status to track the UFD group status.
VLT port-channel configuration removed from the port-channel interface which is upstream member of UFD group	No action	No action	Stops tracking the fabric status for the UFD group. Starts tracking the local port-channel operational status, which is upstream member of the UFD group.
Fabric Status is operationally up	No action	No action	Enables the downstream members, that is clears the error-disabled state.
Fabric Status is operationally down	No action	No action	Disables the downstream members, that is sets the error-disabled state.

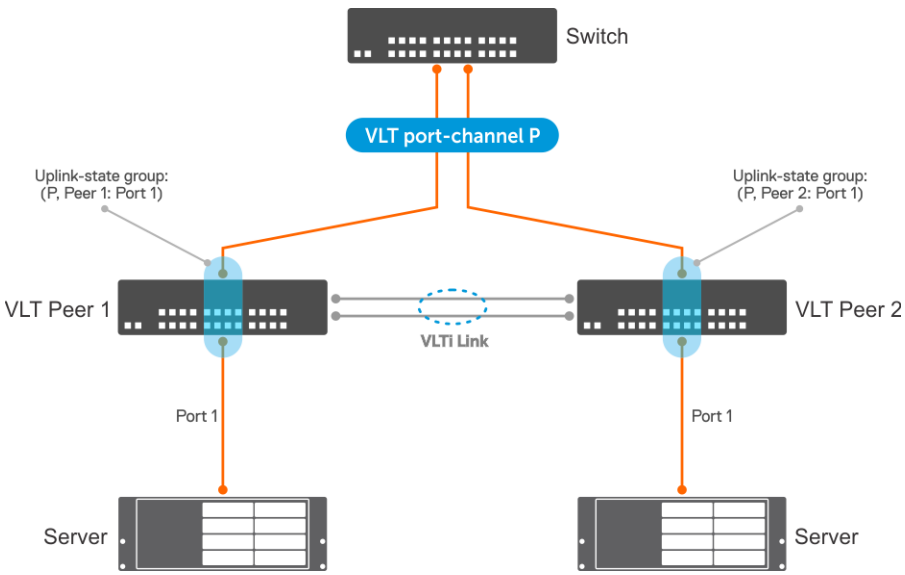
Sample configurations of UFD on VLT

The following examples show some of the uplink-state groups on VLT.

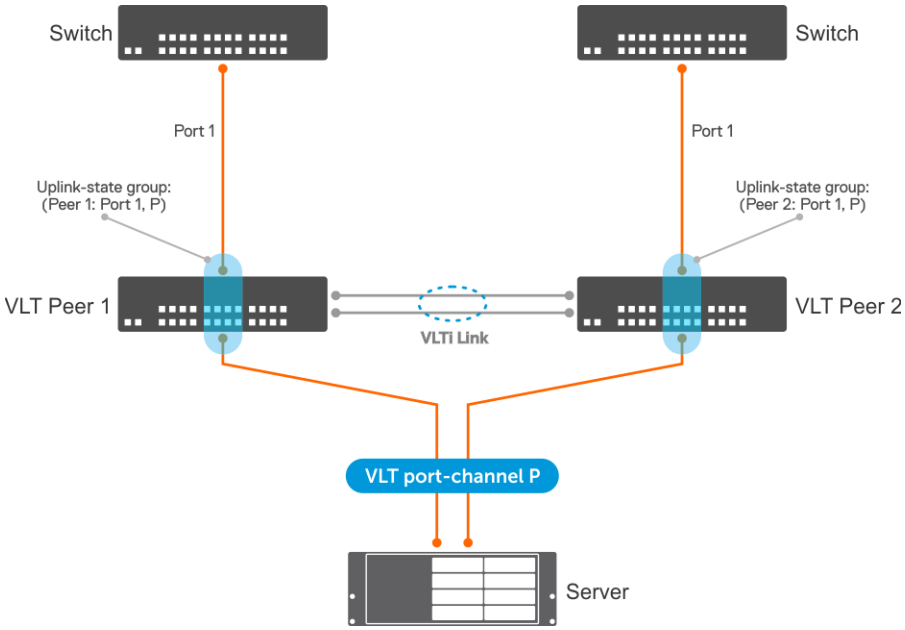
In the following illustration, both the upstream and downstream members are part of VLT port-channels. The uplink-state group includes both the VLT port-channels as members.



In the following example, the upstream member is part of VLT port-channel and the downstream member is an orphan port. The uplink-state group includes the VLT port-channel, VLT node, and the downstream port. The configuration is symmetric on both the VLT nodes.



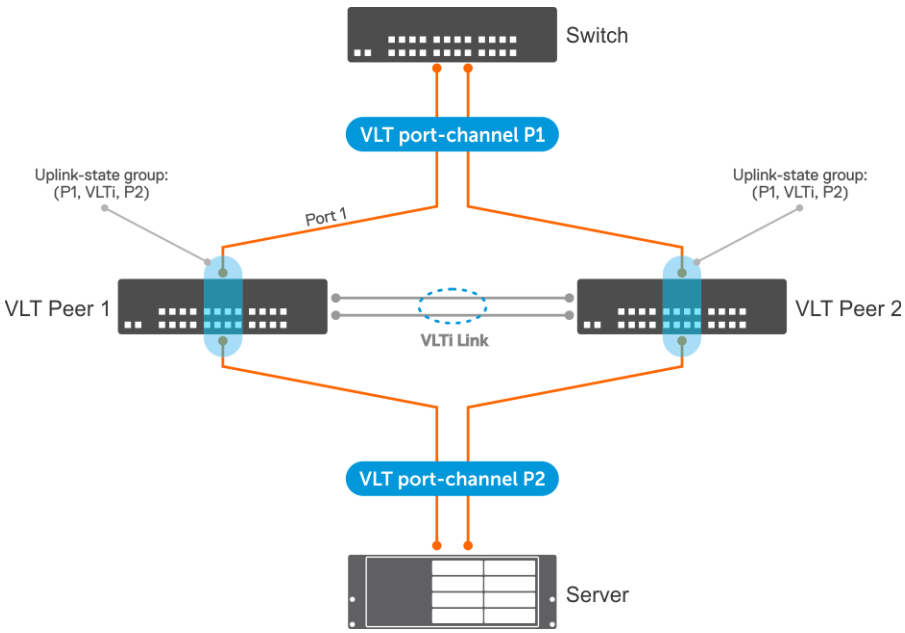
In the following example, the downstream member is part of VLT port-channel and the upstream member is an orphan port. The uplink-state group includes the VLT port-channel, VLT node, and the upstream port. The configuration is symmetric on both the VLT nodes.



OS10 does not support adding a VLTi link member to the uplink-state group. You can add the VLTi link as upstream member to an uplink-state group using the `upstream VLTi` command. If the VLTi link is not available in the system, OS10 allows adding the VLTi link as an upstream member. In this case, UFD starts tracking the operational status of the VLTi link when the link is available. Until the VLTi link is available, the `show uplink-state-group details` command displays the status of the link as `NA`.

In the following example, both the VLT port-channel connected to the switch and the VLTi Link are upstream members. The VLT port-channel connected to the server is a downstream member. The UFD tracks the operational status of the peer port-channel.

NOTE: You cannot configure a VLTi link as a downstream member in an uplink-state group. If you configure, UFD disables the VLTi link when the upstream members are operationally down, which affects the VLT functionality.



UFD commands

clear ufd-disable

Overrides the uplink-state group configuration and brings up the downstream interfaces.

Syntax `clear ufd-disable {interface interface-type | uplink-state-group group-id}`

Parameters

- *interface-type* — Enter the interface type.
- *group-id* — Enter the uplink state group ID, from 1 to 32.

Default None

Command Mode EXEC

Usage Information This command manually brings up a disabled downstream interface that is in an UFD-disabled error state. After the downstream interface is up, it is not disabled until there are changes in the upstream interfaces. This command does not affect downstream interfaces that are already up or interfaces that are not part of the UFD group.

Example

```
OS10# clear ufd-disable interface ethernet 1/1/2
OS10# clear ufd-disable uplink-state-group 1
```

Supported Releases 10.4.0E(R3) or later

defer-time

Configures the timer to defer UFD actions on downstream ports.

Syntax `defer-time timer`

Parameters *timer*— Enter the timer value in seconds, ranging from 1 to 120.

Default Disabled

Command Mode UPLINK-STATE-GROUP

Usage Information You can view configured timer details using the `show uplink-state-group [group-id] detail` command. The `no` version of this command disables the timer.

Example

```
OS10(config)# uplink-state-group 1
OS10(conf-uplink-state-group-1)# defer-time 120
```

Supported Releases 10.4.1.0 or later

downstream

Adds an interface or a range of interfaces as a downstream interface to the uplink-state group.

Syntax `downstream {interface-type | interface-range}`

Parameters

- *interface-type* — Enter the interface type as Ethernet or port-channel.
- *interface-range* — Enter the range of interfaces.

Default None

Command Mode	UPLINK-STATE-GROUP
Usage Information	You cannot assign an interface that is already a member of an uplink-state group to another group. The `no` version of this command removes the interface from the uplink-state group.
Example	`OS10(config)# uplink-state-group 1` `OS10(conf-uplink-state-group-1)# downstream ethernet 1/1/1`
Supported Releases	10.4.0E(R3) or later

downstream auto-recover

Enables auto-recovery of the disabled downstream interfaces.

Syntax	`downstream auto-recover`
Parameters	None
Default	Enabled
Command Mode	UPLINK-STATE-GROUP
Usage Information	The `no` version of this command disables the auto-recovery of downstream interfaces.
Example	`OS10(config)# uplink-state-group 1` `OS10(conf-uplink-state-group-1)# no downstream auto-recover`
Supported Releases	10.4.1.0 or later

downstream disable links

Configures the number of downstream interfaces to disable when an upstream interface in the uplink-state group goes down.

Syntax	`downstream disable links {`*`number`*` | all}`
Parameters	• *`number`*—Enter the number of downstream interfaces to disable, from 1 to 1024. • `all`—Enter `all` to disable all the downstream interfaces.
Default	Not configured
Command Mode	UPLINK-STATE-GROUP
Usage Information	The `no` version of this command reverts the settings to the default state.
Example	`OS10(config)# uplink-state-group 1` `OS10(conf-uplink-state-group-1)# downstream disable links 2`
Supported Releases	10.4.1.0 or later

enable

Enables tracking of an uplink-state group.

Syntax	`enable`

Parameters	None
Default	Disabled
Command Mode	UPLINK-STATE-GROUP
Usage Information	The `no` version of this command disables tracking of an uplink-state group.
Example	`OS10(config)# uplink-state-group 1` `OS10(conf-uplink-state-group-1)# enable`
Supported Releases	10.4.0E(R3) or later

name

Configures a descriptive name for the uplink-state group.

Syntax	`name` *`string`*
Parameters	*`string`* — Enter a description for the uplink-state group. A maximum of 32 characters.
Default	Not configured
Command Mode	UPLINK-STATE-GROUP
Usage Information	The `no` version of this command removes the descriptive name.
Example	`OS10(config)# uplink-state-group 1` `OS10(conf-uplink-state-group-1)# name test_ufd_group`
Supported Releases	10.4.0E(R3) or later

show running-configuration uplink-state-group

Displays the running configuration specific to uplink-state groups.

Syntax	`show running-configuration uplink-state-group [`*`group-id`*`]`
Parameters	*`group-id`* — Enter the uplink group ID. The running configuration of the specified group ID displays.
Default	Not configured
Command Mode	EXEC
Usage Information	None

Example	`OS10# show running-configuration uplink-state-group`
	`!`
	`uplink-state-group 1`
	` downstream ethernet1/1/8:1-1/1/8:4`
	` upstream ethernet1/1/9:1-1/1/9:4`
	` upstream port-channel1-3`

Supported Releases	10.4.0E(R3) or later

show uplink-state-group

Displays the configured uplink-state status.

Syntax

```
show uplink-state-group [group-id] [detail]
```

Parameters

- *group-id* — Enter the uplink group ID. The status of the specified group ID displays.
- detail — Displays detailed information on the status of the uplink-state groups.

Default

Not configured

Command Mode

EXEC

Usage Information

None

Example

```
OS10# show uplink-state-group

Uplink State Group: 9, Status: Enabled,down
```

```
OS10# show uplink-state-group 9

Uplink State Group: 9, Status: Enabled,down
OS10#
```

Example (detail)

```
OS10# show uplink-state-group detail

(Up): Interface up   (Dwn): Interface down   (Dis): Interface disabled

Uplink State Group   :  1     Status   :  Enabled,up Name :  UFDGROUP1
Defer Time           :  10 second(s)
Upstream Interfaces  :  Eth 1/1/7:1(Up)
Downstream Interfaces:  Eth 1/1/1(Dwn)   Eth 1/1/2(Dwn)   Eth 1/1/3(Dwn)   Eth
1/1/4(Dwn)
                        Eth 1/1/5(Dwn)   Eth 1/1/9:2(Dwn) Eth 1/1/9:3(Dwn)
```

```
OS10# show uplink-state-group 2 detail

(Up): Interface up   (Dwn): Interface down   (Dis): Interface disabled

Uplink State Group   :  2      Status  :  Enabled,down  Name:  UFDGROUP
Upstream Interfaces  :  Eth 1/1/6(Dwn)   Eth 1/1/10(Dwn)  Eth 1/1/11(Dwn)  Eth
1/1/12(Dwn)
                        Eth 1/1/13(Dwn)  Eth 1/1/14(Dwn)  Eth 1/1/15(Dwn)
Downstream Interfaces:  Eth 1/1/16(Dis)  Eth 1/1/17(Dis)  Eth 1/1/18(Dis)  Eth
1/1/19(Dis)
                        Eth 1/1/20(Dis)
```

Example (detail with VLTi and VLT status tracked)

```
OS10#show uplink-state-group 1 detail

(Up): Interface up (Dwn): Interface down (Dis): Interface disabled (NA): Not
Available

*: VLT port-channel, V: VLT status, P: Peer Operational status ^: Tracking
status

Uplink State Group : 1 Name: iscsi_group, Status: Enabled, Up
Upstream Interfaces : eth1/1/35(Up) *po10(V:Up, ^P:Dwn) VLTi(NA)
Downstream Interfaces : eth1/1/2(Up) *po20(V: Up,P: Up)
```

```
OS10#show uplink-state-group 2 detail

(Up): Interface up (Dwn): Interface down (Dis): Interface disabled (NA): Not
```

```
Available

*: VLT port-channel, V: VLT status, P: Peer Operational status ^: Tracking
status

Uplink State Group : 1 Name: iscsi_group, Status: Enabled, Up
Upstream Interfaces : eth1/1/36(Up) *po30(^V:Up, P:Dwn) VLTi(Up)
Downstream Interfaces : eth1/1/4(Up) *po20(V: Up,P: Up)
```

Supported Releases 10.4.0E(R3) or later

uplink-state-group

Creates an uplink-state group and enables upstream link tracking.

Syntax `uplink-state-group group-id`

Parameters *group-id* — Enter a unique ID for the uplink-state group, from 1 to 32.

Default None

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the uplink-state group.

Example

```
OS10(config)# uplink-state-group 1
```

Supported Releases 10.4.0E(R3) or later

upstream

Adds an interface or a range of interfaces as an upstream interface to the uplink-state group.

Syntax `upstream {interface-type | interface-range [ track-vlt-status ] | VLTi}`

Parameters

- *interface-type* — Enter the interface type as Ethernet or port-channel.
- *interface-range* — Enter the range of interfaces.
- `VLTi`—Configures VLTi Link as member of uplink state group.
- `track-vlt-status`—(Optional) Tracks the VLT status for the upstream member. This option applies only for port-channel interfaces.

Default When you add an upstream member without the `track-vlt-status` option, the operational status is tracked by default.

Command Mode UPLINK-STATE-GROUP

Usage Information You cannot assign an interface that is already a member of an uplink-state group to another group. The `no` version of this command removes the interface from the uplink-state group.

Example

```
OS10(config)# uplink-state-group 1
OS10(conf-uplink-state-group-1)# upstream ethernet 1/1/45-1/1/48
```

```
OS10(conf-uplink-state-group-1)# upstream VLTi
```

```
OS10(conf-uplink-state-group-1)# upstream port-channel 10 track-vlt-status
```

Supported Releases 10.4.0E(R3) or later

15

Converged data center services

OS10 supports converged data center services, including IEEE 802.1 data center bridging (DCB) extensions to classic Ethernet. DCB provides I/O consolidation in a data center network. Each network device carries multiple traffic classes while ensuring lossless delivery of storage traffic with best-effort for local area network (LAN) traffic and latency-sensitive scheduling of service traffic.

- 802.1Qbb — Priority flow control
- 802.1Qaz — Enhanced transmission selection
- Data Center Bridging Exchange (DCBX) protocol

DCB enables the convergence of LAN and storage area network (SAN) traffic over a shared physical network in end-to-end links from servers to storage devices. In a converged network, all server, storage, and networking devices are DCB-enabled. DCB supports fibre channel over Ethernet (FCoE) and iSCSI transmission of storage data. DCB is not supported on interfaces with link-level flow control (LLFC) enabled.

Priority flow control (PFC)	Use priority-based flow control to ensure lossless transmission of storage traffic, while transmitting other traffic classes that perform better without flow control, see Priority flow control.
Enhanced transmission selection (ETS)	Assign bandwidth to 802.1p class of service (CoS)-based traffic classes. Use ETS to increase preferred traffic-class throughput during network congestion, see Enhanced transmission selection.
Data Center Bridging Exchange protocol (DCBX)	Configure the DCBX protocol DCB neighbors use to discover and exchange configuration information for plug-and-play capability, see Data center bridging eXchange.
Internet small computer system interface (iSCSI)	Use iSCSI auto-configuration and detection of storage devices, monitor iSCSI sessions, and apply QoS policies on iSCSI traffic, see Internet small computer system interface.

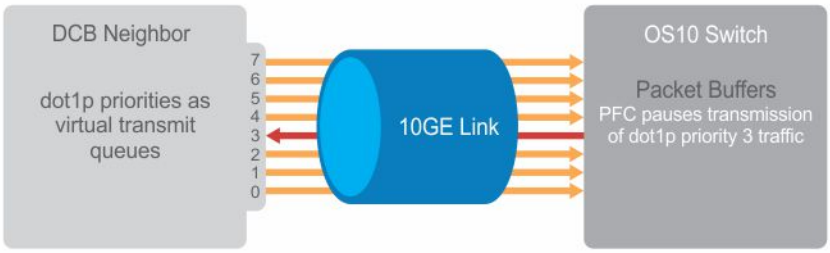
Priority flow control

In a converged data-center network, to ensure that no frames are lost due to congestion, use PFC. PFC uses the 802.1p priority in the Ethernet header to pause priority-specific traffic sent from a transmitting device. The 802.1p priority is also known as the class of service (CoS) or dot1p priority value.

When PFC detects congestion of a dot1p traffic class, it sends a pause frame for the priority traffic to the transmitting device. In this way, PFC ensures that the switch does not drop specified priority traffic.

PFC enhances the existing 802.3x pause capability to enable flow control based on 802.1p priorities. Instead of stopping all traffic on a link, as performed by the 802.3x pause mechanism, PFC pauses traffic for 802.1p traffic types. For example, when LAN traffic congestion occurs on an interface, PFC ensures lossless flows of storage and server traffic while allowing for lossy best-effort transmission of other traffic.

PFC handles traffic congestion by pausing prioritized dot1p traffic on an ingress interface and allowing other dot1p traffic best-effort, also known as lossy data transmission.



PFC configuration notes

- PFC is supported for 802.1p, dot1p priority traffic, from 0 to 7. FCoE traffic traditionally uses dot1p priority 3 — iSCSI storage traffic uses dot1p priority 4.
- Configure PFC for ingress traffic by using network-qos class and policy maps, see *Quality of Service*. PFC-enabled traffic queues are treated as lossless queues. Configure the same network-qos policy map on all PFC-enabled ports. Configure required bandwidth for lossless traffic using ETS queuing (output) policies on egress interfaces.
- In a network-qos policy-class map, use commands to generate PFC pause frames for matching class-map priorities:
  - Send pause frames for matching class-map traffic during congestion using the `pause` command.
  - (Optional) Enter user-defined values for the reserved ingress buffer-size of PFC class-map traffic, and the thresholds used to send XOFF and XON pause frames using the `pause [buffer-size` *`kilobytes`* `pause-threshold` *`kilobytes`* `resume-threshold` *`kilobytes`*`]` command.
  - Configure the matching dot1p values used to send pause frames using the `pfc-cos` command.
  - (Optional) Set the static and dynamic thresholds that determine the shared buffers available for PFC class-map traffic queues using the `queue-limit thresh-mode` command.
- By default, all ingress traffic is handled by the lossy ingress buffer. When you enable PFC, dot1p ingress traffic competes for shared buffers in the lossless pool instead of the shared lossy pool. The number of lossless queues supported on an interface depends on the amount of available free memory in the lossy pool.
- Use the `priority-flow-control mode on` command to enable PFC for FCoE and iSCSI traffic; for example, priority 3 and 4.
- Enable DCBX on interfaces to detect and auto-configure PFC/ETS parameters from peers.
- PFC and 802.3x LLFC are disabled by default on an interface. You cannot enable PFC and LLFC at the same time. LLFC ensures lossy traffic in best-effort transmission. Enable PFC to enable guarantee lossless FCoE and iSCSI traffic. PFC manages buffer congestion by pausing specified ingress dot1p traffic; LLFC pauses all data transmission on an interface. To enable LLFC, use the `flowcontrol [receive | transmit] [on | off]` command.
- SYSTEM-QOS mode applies a service policy globally on all interfaces:
  - Create and apply a 1-to-1 802.1p-priority-to-traffic-class mapping on an interface or all interfaces in INTERFACE or SYSTEM-QOS mode
  - Create and apply a 1-to-1 traffic-class-to-queue mapping on an interface or all interfaces in INTERFACE or SYSTEM-QOS mode

The S5148F-ON platform has the following limitations:

- You cannot configure PFC priority 0 as a lossless priority.
- You cannot map multiple priorities to the same queue.
- Whenever you enable LLFC on an interface, Rx PFC frames are honored. Also, whenever you enable PFC on an interface, Rx Pause frames are honored. Rx Pause statistics in the hardware also includes the Rx PFC frames.

Configure dot1p priority to traffic class mapping

Decide if you want to use the default 802.1p priority-to-traffic class (`qos-group`) mapping or configure a new map. By default, the qos class-trust class map is applies to ingress traffic. The class-trust class instructs OS10 interfaces to honor dot1p or differentiated services code point (DSCP) traffic.

```
 Dot1p Priority : 0  1  2  3  4  5  6  7
  Traffic Class : 1  0  2  3  4  5  6  7
```

- Apply the default trust map specifying that dot1p values are trusted in SYSTEM-QOS or INTERFACE mode.

  ```
  trust-map dot1p default
  ```

Configure a non-default dot1p-priority-to-traffic class mapping

1 Configure a trust map of dot1p traffic classes in CONFIGURATION mode. A trust map does not modify ingress dot1p values in output flows.

  Assign a `qos-group` to trusted dot1p values in TRUST mode using 1-to-1 mappings. Dot1p priorities are 0 to 7. For a PFC traffic class, map only one `dot1p` value to a `qos-group` number; for Broadcom-based NPU platforms, the `qos-group` number and the `dot1p` value must be the same. A `qos-group` number is used only internally to classify ingress traffic classes.

  ```
  trust dot1p-map dot1p-map-name
     qos-group {0-7} dot1p {0-7}
     exit
  ```

2 Apply the `trust dot1p-map` policy to ingress traffic in SYSTEM-QOS or INTERFACE mode.

  ```
  trust-map dot1p trust-policy—map-name
  ```

Configure traffic-class-queue mapping

Decide if you want to use the default traffic-class-queue mapping or configure a non-default traffic-class-to-queue mapping.

```
 Traffic Class : 0  1  2  3  4  5  6  7
         Queue : 0  1  2  3  4  5  6  7
```

If you are using the default traffic-class-to-queue map, no further configuration steps are necessary.

1 Create a traffic-class-to-queue map in CONFIGURATION mode. Assign a traffic class (`qos-group`) to a queue in QOS-MAP mode using 1-to-1 mappings. For a PFC traffic class, map only one `qos-group` value to a `queue` number. A `qos-group` number is used only internally to classify ingress traffic.

  ```
  qos—map traffic-class tc-queue-map-name
     queue {0-7} qos-group {0-7}
     exit
  ```

2 Apply the traffic-class-queue map in SYSTEM-QOS or INTERFACE mode.

  ```
  qos-map traffic-class tc-queue-map-name
  ```

View interface PFC configuration

```
OS10# show interface ethernet 1/1/1 priority-flow-control details
ethernet1/1/1
Admin Mode : true
Operstatus: true
PFC Priorities: 4
Total Rx PFC Frames: 0
Total Tx PFC frames: 0
Cos    Rx               Tx
---------------------------------------
0      0              0
1      0              0
2      0              0
3      0              0
4      0              0
5      0              0
6      0              0
7      0              0
```

Configure PFC

PFC provides a pause mechanism based on the 802.1p priorities in ingress traffic. PFC prevents frame loss due to network congestion. Configure PFC lossless buffers, and enable pause frames for dot1p traffic on a per-interface basis. Repeat the PFC configuration on each PFC-enabled interface. PFC is disabled by default.

Decide if you want to use the default dot1p-priority-to-traffic class mapping and the default traffic-class-to-queue mapping. To change the default settings, see PFC configuration notes.

Configuration steps:

1 Create PFC dot1p traffic classes.
2 Configure ingress buffers for PFC traffic.
3 Apply a service policy and enable PFC.
4 (Optional) Configure the PFC shared buffer for lossless traffic.

Create PFC dot1p traffic classes

1 Create a `network-qos` class map to classify PFC traffic classes in CONFIGURATION mode, from 1 to 7. Specify the traffic classes using the `match qos-group` command. QoS-groups map 1:1 to traffic classes 1 to 7; for example, `qos-group 1` corresponds to traffic class 1. Enter a single value, a hyphen-separated range, or multiple `qos-group` values separated by commas in CLASS-MAP mode.

```
class–map type network-qos class–map-name
   match qos-group {1-7}
   exit
```

2 (Optional) Repeat Step 1 to configure additional PFC traffic-class class-maps.

NOTE: In the S5148F-ON, PFC is not supported on priority 0.

Configure pause and ingress buffers for PFC traffic

For the default ingress queue settings and the default dot1p priority-queue mapping, see PFC configuration notes.

1 Create a `network-qos` policy map in CONFIGURATION mode.

```
policy-map type network-qos policy–map-name
```

2 Associate the policy-map with a `network-qos` class map in POLICY-MAP mode.

```
class class-map-name
```

3 Configure default values for ingress buffers used for the `network-qos` class maps in POLICY-CLASS-MAP mode.

```
pause
```

(Optional) Change the default values for the ingress-buffer size reserved for the `network-qos` class-map traffic and the thresholds used to send XOFF and XON pause frames in kilobytes.

```
pause [buffer-size kilobytes {pause-threshold kilobytes | resume-threshold kilobytes}]
```

4 Enable the PFC pause function for dot1p traffic in POLICY-CLASS-MAP mode. The dot1p values must be the same as the `qos-group` traffic class numbers in the class map in Step 2. Enter a single dot1p value, from 1 to 7, a hyphen-separated range, or multiple dot1p values separated by commas.

```
pfc-cos dot1p-priority
```

5 (Optional) Set the static and dynamic thresholds used to limit the shared buffers allocated to PFC traffic-class queues. Configure a static, fixed queue-limit (in kilobytes) or a dynamic threshold (weight 1-10; default 9) based on the available PFC shared buffers. This option is not available in S5148F-ON.

```
queue-limit thresh-mode {static kilobytes | dynamic weight}
```

6 (Optional) Repeat Steps 2–4 to configure PFC on additional traffic classes.

Apply service policy and enable PFC

1 Apply the PFC service policy on an ingress interface or interface range in INTERFACE mode.

```
interface ethernet node/slot/port:[subport]
   service-policy input type network-qos policy—map-name
```

```
interface range ethernet node/slot/port:[subport]-node/slot/port[:subport]
   service-policy input type network-qos policy—map-name
```

2 Enable PFC without DCBX for FCoE and iSCSI traffic in INTERFACE mode.

```
priority-flow-control mode on
```

Configure PFC

PFC is enabled on traffic classes with dot1p 3 and 4 traffic. The two traffic classes require different ingress queue processing. In the `network-qos pp1` policy map, class `cc1` uses customized PFC buffer size and pause frame settings; class `cc2` uses the default settings.

```
OS10(config)# system qos
OS10(config-sys-qos)# trust-map dot1p default

OS10(config)# system qos
OS10(config-sys-qos)# service-policy input type qos pclass1
OS10(config-sys-qos)# exit

OS10(config)# class-map type network-qos cc1
OS10(config-cmap-nqos)# match qos-group 3
OS10(config-cmap-nqos)# exit

OS10(config)# class-map type network-qos cc2
OS10(config-cmap-nqos)# match qos-group 4
OS10(config-cmap-nqos)# exit

OS10(config)# policy-map type network-qos pp1
OS10(config-pmap-network-qos)# class cc1
OS10(config-pmap-c-nqos)# pause buffer-size 30 pause-threshold 20 resume-threshold 10
OS10(config-pmap-c-nqos)#pfc-cos 3
OS10(config-pmap-c-nqos)#exit
OS10(config-pmap-network-qos)# class cc2
OS10(config-pmap-c-nqos)# pause
OS10(config-pmap-c-nqos)#pfc-cos 4
OS10(config-pmap-c-nqos)#exit

OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# service-policy input type network-qos pp1

OS10(conf-if-eth1/1/1)# priority-flow-control mode on
OS10(conf-if-eth1/1/1)# no shutdown
```

View PFC configuration and operational status

```
OS10(conf-if-eth1/1/1)# do show interface ethernet 1/1/1 priority-flow-control details
ethernet1/1/1
Admin Mode : true
Operstatus: true
PFC Priorities: 3,4
Total Rx PFC Frames: 300
Total Tx PFC frames: 200
Cos    Rx             Tx
-------------------------
0      0              0
1      0              0
2      0              0
3      300            200
4      0              0
5      0              0
6      0              0
7      0              0
```

View PFC ingress buffer configuration

```
OS10# show qos ingress buffers interface ethernet 1/1/1
Interface : ethernet1/1/1
Speed : 0
Priority-grp      Reserved        Shared-buffer       Shared-buffer
XOFF                                                                    XON
    no           buffer-size         mode              threshold
threshol                                                              d
threshold
-------------------------------------------------------------------------------

-----------------------------------------
0                 -                 -                 -
-                                                                       -
1                 -                 -                 -
-                                                                       -
2                 -                 -                 -
-                                                                       -
3                 -                 -                 -
-                                                                       -
4                 -                 -                 -
-                                                                       -
5                 -                 -                 -
-                                                                       -
6                 -                 -                 -
-                                                                       -
7                 9360              static            12779520
-                                                                       -
```

View PFC system buffer configuration

```
OS10# show qos system ingress buffer
All values are in kb
Total buffers                                - 12187
    Total lossless buffers                   - 0
      Maximum lossless buffers               - 5512
      Total shared lossless buffers          - 0
      Total used shared lossless buffers     -
    Total lossy buffers                      - 11567
      Total shared lossy buffers             - 11192
      Total used shared lossy buffers        - 0


OS10# show qos system egress buffer
All values are in kb
Total buffers                              - 12187
    Total lossless buffers                 - 0
      Total shared lossless buffers        - 0
      Total used shared lossless buffers   -
    Total lossy buffers                    - 11567
      Total shared lossy buffers           - 9812
      Total used shared lossy buffers      - 0
    Total CPU buffers                    - 620
      Total shared CPU buffers           - 558
      Total used shared CPU buffers        - 0
```

View PFC ingress buffer statistics

```
OS10(config)# show qos ingress buffer-stats interface ethernet 1/1/15
Interface : ethernet1/1/15
Speed : 10G
Priority Used reserved     Used shared  Used HDRM
 Group   buffers           buffers      buffers
-----------------------------------------------
0         9360              681824       35984
1         0                 0            0
2         0                 0            0
3         0                 0            0
```

```
4          0                    0              0
5          0                    0              0
6          0                    0              0
7          0                    0              0
```

PFC commands

pause

Configures the ingress buffer and pause frame settings used for PFC traffic classes.

Syntax

`pause [buffer-size` *`kilobytes`* `pause-threshold` *`kilobytes`* `resume-threshold` *`kilobytes`*`]`

Parameters

- `buffer-size` *`kilobytes`* — Enter the reserved (guaranteed) ingress-buffer size in kilobytes for PFC dot1p traffic, from 0 to 7787.
- `pause-threshold` *`kilobytes`* — Enter the threshold used to send pause frames in kilobytes to a transmitting device, from 0 to 7787.
- `resume-threshold` *`kilobytes`* — Enter the threshold used to request a transmitting device in kilobytes to resume sending traffic, from 0 to 7787.

Defaults

The default ingress-buffer size reserved for PFC traffic classes, and the pause and resume thresholds, vary according to the interface type. The default egress buffer reserved for PFC traffic classes is 0 on all interface types.

Table 70. Port defaults

Port Speed	10G Port	25G Port	40G Port	100G Port
PFC reserved ingress buffer	166 KB	195 KB	315.5 KB	512 KB
PFC pause threshold	96 KB	96 KB	192 KB	232 KB
PFC resume threshold	87 KB	87 KB	183 KB	223 KB

Command Mode

POLICY-CLASS NETWORK-QOS

Usage Information

Use the `pause` command without optional parameters to apply the default ingress-buffer size, and pause (XON) and resume (XOFF) thresholds. Default values for the `buffer-size`, `pause-threshold` and `resume-threshold` parameters vary across interface types and port speeds. The default values are based on the default MTU size of 9216 bytes.

Example

```
OS10(config)# policy-map type network-qos pp1
OS10(conf-pmap-network-qos)# class cc1
OS10(conf-pmap-c-nqos)# pause buffer-size 30 pause-threshold 20 resume-
threshold 10
```

Supported Releases

10.3.0E or later

pfc-cos

Configures the matching dot1p values used to send PFC pause frames.

Syntax `pfc-cos dot1p-priority`

Parameters *dot1p-priority* — Enter a single dot1p priority value for a PFC traffic class, from 1 to 7, a hyphen-separated range, or multiple dot1p values separated by commas.

Default Not configured

Command Mode POLICY-CLASS NETWORK-QOS

Usage Information When you enter PFC-enabled dot1p priorities with `pfc-cos`, the dot1p values must be the same as the `match qos-group` (traffic class) numbers in the network-qos class map used to define the PFC traffic class, see *Configure PFC Example*. A `qos-group` number is used only internally to classify ingress traffic classes. For the default dot1p-priority-to-traffic-class mapping and how to configure a non-default mapping, see PFC configuration notes. A PFC traffic class requires a 1-to-1 mapping — only one `dot1p` value is mapped to a `qos-group` number.

Example

```
OS10(config)# class-map type network-qos cc1
OS10(conf-cmap-nqos)# match qos-group 3
OS10(conf-cmap-nqos)# exit
```

Example (policy-map)

```
OS10(config)# policy-map type network-qos pp1
OS10(conf-pmap-network-qos)# class cc1
OS10(conf-pmap-c-nqos)# pfc-cos 3
```

Supported Releases 10.3.0E or later

pfc-shared-buffer-size

Configures the amount of shared buffers available for PFC-enabled traffic on the switch.

Syntax `pfc-shared-buffer-size kilobytes`

Parameter *kilobytes* — Enter the total amount of shared buffers available to PFC-enabled dot1p traffic in kilobytes, from 0 to 7787.

Default 832KB

Command Mode SYSTEM-QOS

Usage Information By default, the lossy ingress buffer handles all ingress traffic. When you enable PFC, dot1p ingress traffic competes for shared buffers in the lossless pool instead of the shared lossy pool. Use this command to increase or decrease the shared buffer allowed for PFC-enabled flows. The configured amount of shared buffers is reserved for PFC flows only after you enable PFC on an interface using the `priority-flow-control mode on` command.

Example

```
OS10(config)# system qos
OS10(conf-sys-qos)# pause-shared-buffer-size 1024
```

Supported Releases 10.3.0E or later

priority-flow-control

Enables PFC on ingress interfaces.

Syntax `priority-flow-control {mode on}`

Parameter `mode on` — Enable PFC for FCoE and iSCSI traffic on an interface without enabling DCBX.

Default Disabled

Command Mode INTERFACE

Usage Information Before you enable PFC, apply a network-qos policy-class map with the specific PFC dot1p priority values to the interface. In the PFC network-qos policy-class map, use the default `buffer-size` values if you are not sure about the `pause-threshold` and `resume-threshold` settings that you want to use. You cannot enable PFC and LLFC at the same time on an interface. The `no` version of this command disables PFC on an interface. When you disable PFC, remove the PFC `network-qos` policy-class map applied to the interface.

Example

```
OS10(conf-if-eth1/1/1)# priority-flow-control mode on
```

Supported Releases 10.3.0E or later

queue-limit

Sets the static and dynamic thresholds used to limit the shared-buffer size of PFC traffic-class queues.

Syntax `queue-limit {thresh-mode [static` *`kilobytes`* `| dynamic` *`weight`*`]}`

Parameters

- `thresh-mode` — Buffer threshold mode.
- `static` *`kilobytes`* — Enter the fixed shared-buffer limit available for PFC traffic-class queues in kilobytes, from 0 to 7787; maximum amount tuned by the `pfc-shared-buffer-size` command.
- `dynamic` *`weight`* — Enter the weight value used to dynamically determine the shared-buffer limit available for PFC traffic-class queuesm from 1 to 10.

Default Dynamic weight of 9 and static shared-buffer limit of 12479488 kilobytes

Command Mode POLICY-CLASS NETWORK-QOS

Usage Information To tune the amount of shared buffers available for the static limit of PFC traffic-class queues on the switch, use the `pfc-shared-buffer-size` command. The current amount of available shared buffers determines the dynamic queue-limit.

Example

```
OS10(config)# policy-map type network-qos pp1
OS10(conf-pmap-network-qos)# class cc1
OS10(conf-pmap-c-nqos)# queue-limit thresh-mode static 1024
```

Supported Releases 10.3.0E or later

show interface priority-flow-control

Displays PFC operational status, configuration, and statistics on an interface.

Syntax `show interface [ethernet` *`node/slot/port`*`[:`*`subport`*`]] priority-flow-control [details]`

Parameters	None
Default	Not configured
Command Mode	EXEC
Usage Information	Use the `details` option to display PFC statistics on received/transmitted frames for each dot1p CoS value. Use the `clear qos statistics interface ethernet 1/1/1` command to delete PFC statistics and restart the counter.

Example (details)

```
OS10(config)# show interface ethernet 1/1/15 priority-flow-control details

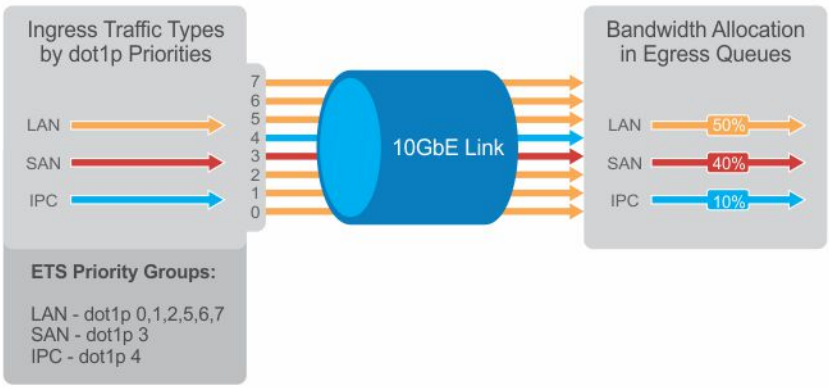
ethernet1/1/15
Admin Mode : true
Operstatus: true
PFC Priorities: 3
Total Rx PFC Frames: 0
Total Tx PFC frames: 587236
Cos    Rx         Tx
----------------------
0      0          0
1      0          0
2      0          0
3      0          587236
4      0          0
5      0          0
6      0          0
7      0          0
```

Supported Releases	10.3.0E or later

Enhanced transmission selection

ETS provides customized bandwidth allocation to 802.1p classes of traffic. Assign different amounts of bandwidth to Ethernet, FCoE, or iSCSI traffic classes that require different bandwidth, latency, and best-effort treatment during network congestion.

ETS divides traffic into different priority groups using their 802.1p priority value. To ensure that each traffic class is correctly prioritized and receives the required bandwidth, configure bandwidth and queue scheduling for each priority group. To prioritize low-latency storage and server-cluster traffic, allocate more bandwidth to a priority group. To rate-limit best-effort LAN traffic, allocate less bandwidth to a different priority group.



ETS configuration notes

- ETS is supported on Layer2 (L2) 802.1p priority (dot1p 0 to 7) and Layer 3 (L3) DSCP (0 to 63) traffic. FCoE traffic uses dot1p priority 3 — iSCSI storage traffic uses dot1p priority 4.
- Apply these maps and policies on interfaces:
  - Trust maps — OS10 interfaces do not honor the L2 and L3 priority fields in ingress traffic by default. Create a trust map to honor dot1p and DSCP classes of lossless traffic. A trust map does not change ingress dot1p and DSCP values in egress flows. In a trust map, assign a `qos-group` traffic class to trusted dot1p/DSCP values. A qos-group number is used only internally to schedule classes of ingress traffic.
  - QoS map — Create a QoS map to assign trusted dot1p and DSCP traffic classes to lossless queues.
  - Ingress trust policy — Configure a service policy to trust dot1p values in ingress traffic.
  - Egress queuing policy — Configure ETS for egress traffic by assigning bandwidth to matching lossless queues in `queuing` class and policy maps.
- Apply both PFC network-qos (input) and ETS queuing (output) policies on an interface to ensure lossless transmission.
- An ETS-enabled interface operates with dynamic weighted round robin (DWRR) or strict priority scheduling.
- OS10 control traffic is sent to control queues, which have a strict priority that is higher than data traffic queues. ETS-allocated bandwidth is not supported on a strict priority queue. A strict priority queue receives bandwidth only from DCBX type, length, values (TLVs).
- The CEE/IEEE2.5 versions of ETS TLVs are supported. ETS configurations are received in a TLV from a peer.

Configure ETS

ETS provides traffic prioritization for lossless storage, latency-sensitive, and best-effort data traffic on the same link.

- Configure classes of dot1p and DSCP traffic and assign them to lossless queues. Use the class-trust class map to honor ingress dot1p and DSCP traffic.
- Allocate guaranteed bandwidth to each lossless queue. An ETS queue can exceed the amount of allocated bandwidth if another queue does not use its share.

ETS is disabled by default on all interfaces.

1 Configure trust maps of dot1p and DSCP values in CONFIGURATION mode. A trust map does not modify ingress values in output flows. Assign a `qos-group`, traffic class from 0 to 7, to trusted dot1p/DSCP values in TRUST mode. A `qos-group` number is used only internally to schedule classes of ingress traffic. Enter multiple `dot1p` and `dscp` values in a hyphenated range or separated by commas.

```
trust dot1p-map dot1p-map-name
   qos-group {0-7} dot1p {0-7}
   exit
trust dscp-map dscp-map-name
   qos-group {0-7} dscp {0-63}
   exit
```

2 Configure a QoS map with trusted traffic-class (`qos-group`) to lossless-queue mapping in CONFIGURATION mode. Assign one or more qos-groups, from 0 to 7, to a specified queue in QOS-MAP mode. Enter multiple `qos-group` values in a hyphenated range or separated by commas. Enter multiple `queue qos-group` entries, if necessary.

```
qos-map traffic-class queue-map-name
   queue {0-7} qos-group {0-7}
   exit
```

3 Apply the default trust map specifying that dot1p and dscp values are trusted in SYSTEM-QOS or INTERFACE mode.

```
trust-map {dot1p | dscp} default
```

4 Create a `queuing` class map for each ETS queue in CONFIGURATION mode. Enter `match queue` criteria in CLASS-MAP mode.

```
class–map type queuing class–map-name
   match queue {0-7}
   exit
```

5 Create a `queuing` policy map in CONFIGURATION mode. Enter POLICY-CLASS-MAP mode and configure the percentage of bandwidth allocated to each traffic class-queue mapping. The sum of all DWRR-allocated bandwidth across ETS queues must be 100%, not including the strict priority queue. Otherwise, QoS automatically adjusts bandwidth percentages so that ETS queues always receive 100% bandwidth. The remaining non-ETS queues receive 1% bandwidth each.

```
policy—map type queuing policy—map-name
   class class—map-name
      bandwidth percent {1-100}
```

(Optional) To configure a queue as strict priority, use the `priority` command. Packets scheduled to a strict priority queue are transmitted before packets in non-priority queues.

```
policy—map type queuing policy—map-name
   class class—map-name
      priority
```

6 Apply the trust maps for dot1p and DSCP values, and the traffic class-queue mapping globally on the switch in SYSTEM-QOS mode or on an interface or interface range in INTERFACE mode.

```
system qos
   trust-map dot1p dot1p-map-name
   trust-map dscp dscp-map-name
   qos-map traffic-class queue-map-name
```

Or

```
interface {ethernet node/slot/port[:subport] | range ethernet node/slot/port[:subport]-node/
slot/port[:subport]}
   trust-map dot1p dot1p-map-name
   trust-map dscp dscp-map-name
   qos-map traffic-class queue-map-name
```

7 Apply the `qos` trust policy to ingress traffic in SYSTEM-QOS or INTERFACE mode.

```
service-policy input type qos trust-policy—map-name
```

8 Apply the `queuing` policy to egress traffic in SYSTEM-QOS or INTERFACE mode.

```
service-policy output type queuing policy—map-name
```

9 Enable ETS globally in SYSTEM-QOS mode or on an interface/interface range in INTERFACE mode.

```
ets mode on
```

Configure ETS

```
OS10(config)# trust dot1p-map dot1p_map1
OS10(config-trust-dot1pmap)# qos-group 0 dot1p 0-3
OS10(config-trust-dot1pmap)# qos-group 1 dot1p 4-7
OS10(config-trust-dot1pmap)# exit

OS10(config)# trust dscp-map dscp_map1
OS10(config-trust-dscpmap)# qos-group 0 dscp 0-31
OS10(config-trust-dscpmap)# qos-group 1 dscp 32-63
OS10(config-trust-dscpmap)# exit

OS10(config)# qos-map traffic-class tc-q-map1
OS10(config-qos-tcmap)# queue 0 qos-group 0
OS10(config-qos-tcmap)# queue 1 qos-group 1
OS10(config-qos-tcmap)# exit

OS10(config)# system qos
OS10(config-sys-qos)# trust-map dot1p default


OS10(config)# class-map type queuing c1
OS10(config-cmap-queuing)# match queue 0
OS10(config-cmap-queuing)# exit
OS10(config)# class-map type queuing c2
OS10(config-cmap-queuing)# match queue 1
OS10(config-cmap-queuing)# exit

OS10(config)# policy-map type queuing p1
OS10(config-pmap-queuing)# class c1
```

```
OS10(config-pmap-queuing)# bandwidth percent 30
OS10(config-pmap-queuing)# exit
OS10(config)# policy-map type queuing p2
OS10(config-pmap-queuing)# class c2
OS10(config-pmap-queuing)# bandwidth percent 70
OS10(config-pmap-queuing)# exit

OS10(config)# system qos
OS10(config-sys-qos)# trust-map dot1p dot1p_map1
OS10(config-sys-qos)# trust-map dscp dscp_map1
OS10(config-sys-qos)# qos-map traffic-class tc-q-map1
OS10(config-sys-qos)# ets mode on
OS10(config-sys-qos)# service-policy input type qos pclass1
OS10(config-sys-qos)# service-policy output type queuing p1
```

View ETS configuration

```
OS10# show qos interface ethernet 1/1/1
Interface
unknown-unicast-storm-control : Disabled
multicast-storm-control : Disabled
broadcast-storm-control : Disabled
flow-control-rx : Disabled
flow-control-tx : Disabled
ets mode : Disabled
Dot1p-tc-mapping : dot1p_map1
Dscp-tc-mapping : dscp_map1
tc-queue-mapping : tc-q-map1
```

View QoS maps: traffic-class to queue mapping

```
OS10# show qos maps
Traffic-Class to Queue Map: tc-q-map1
      queue 0 qos-group 0
      queue 1 qos-group 1
Traffic-Class to Queue Map: dot1p_map1
      qos-group 0 dot1p 0-3
      qos-group 1 dot1p 4-7
DSCP Priority to Traffic-Class Map  : dscp_map1
      qos-group 0 dscp 0-31
      qos-group 1 dscp 32-63
```

ETS commands

ets mode on

Enables ETS on an interface.

Syntax	`ets mode on`
Parameter	None
Default	Disabled
Command Mode	INTERFACE
Usage Information	Enable ETS on all switch interfaces in SYSTEM-QOS mode or on an interface or interface range in INTERFACE mode. The `no` version of this command disables ETS.
Example	`OS10(config-sys-qos)# ets mode on`
Supported Releases	10.3.0E or later

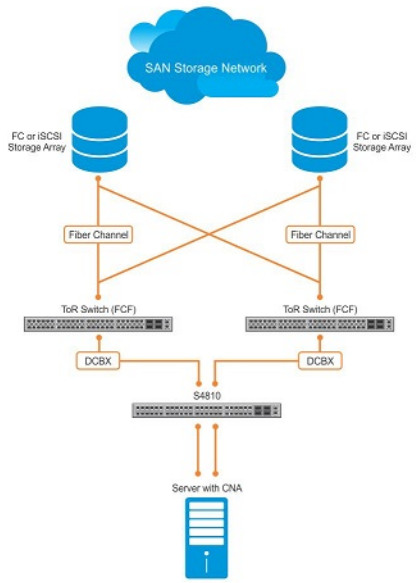
Data center bridging eXchange

DCBX allows a switch to automatically discover and set up DCBX-enabled peers configured with compatible settings. In a converged data center network, DCBX provides plug-and-play capability for server, storage, and networking devices in an end-to-end solution.

DCBX uses link layer discovery protocol (LLDP) to mediate automatic negotiation and device settings exchange, such as PFC and ETS. DCBX uses LLDP TLVs to perform DCB parameter exchange:

- PFC configuration and application-priority configuration
- ETS configuration and ETS recommendation

This sample DCBX topology shows two 40GbE ports on a switch that are configured as DCBX auto-upstream ports and used as uplinks to top-of-rack (ToR) switches. The ToR switches are part of a fibre channel storage network.



DCBX configuration notes

- To exchange link-level configurations in a converged network, DCBX is a prerequisite for using DCB features, such as PFC and ETS. DCBX is also deployed in topologies that support lossless operation for FCoE or iSCSI traffic. In these scenarios, all network devices must be DCBX-enabled so that DCBX is enabled end-to-end.
- DCBX uses LLDP to advertise and automatically negotiate the administrative state and PFC/ETS configuration with directly connected DCB peers. If you disable LLDP on an interface, DCBX cannot run. Enable LLDP on all DCBX ports,
- DCBX is disabled at a global level by default. Enable DCBX globally on a switch to activate the exchange of DCBX TLV messages with PFC, ETS, and iSCSI configurations.
- DCBX is enabled by default on OS10 interfaces. You can manually reconfigure DCBX settings on a per-interface basis. For example, you can disable DCBX on an interface using the `no lldp tlv-select dcbxp` command or change the DCBX version using the `dcbx version` command.
- For DCBX to be operational, DCBX must be enabled at both the global and interface levels. If the `show lldp dcbx interface` command returns the message `DCBX feature not enabled`, DCBX is not enabled at both levels.
- OS10 supports DCBX versions CEE and IEEE2.5.
- By default, DCBX advertises all TLVs—PFC, ETS Recommendation, ETS Configuration, DCBXP, and basic TLVs.

- A DCBX-enabled port operates in a manual role by default. The port operates only with user-configured settings and does not auto-configure with DCB settings received from a DCBX peer. When you enable DCBX, the port advertises its PFC and ETS configurations to peer devices but does not accept external, or propagate internal, DCB configurations.
- DCBX detects misconfiguration on a peer device when DCB features are not compatibly configured with the local switch. Misconfiguration detection is feature-specific because some DCB features support asymmetric (non-identical) configurations.

Configure DCBX

DCBX allows data center devices to advertise and exchange configuration settings with directly connected peers using LLDP. LLDP is enabled by default.

To ensure the consistent and efficient operation of a converged data center network, DCBX detects peer misconfiguration.

DCBX is disabled at a global level and enabled at an interface level by default. For DCBX to be operational, DCBX must be enabled at both the global and interface levels. You can manually reconfigure DCBX settings or disable DCBX on a per-interface basis.

1 Configure the DCBX version used on a port in INTERFACE mode.

```
dcbx version {auto | cee | ieee}
```

   - `auto` — Automatically selects the DCBX version based on the peer response, the default.
   - `cee` — Sets the DCBX version to CEE.
   - `ieee` — Sets the DCBX version to IEEE 802.1Qaz.

2 (Optional) A DCBX-enabled port advertises all TLVs by default. If PFC or ETS TLVs are disabled, enter the command in INTERFACE mode to re-enable PFC or ETS TLV advertisements.

```
dcbx tlv-select {ets-conf | ets-reco | pfc}
```

   - `ets-conf` — Enables ETS configuration TLVs.
   - `ets-reco` — Enables ETS recommendation TLVs.
   - `pfc` — Enables PFC TLVs.

3 (Optional) DCBX is enabled on a port by default. If DCBX is disabled, enable it in INTERFACE mode.

```
lldp tlv-select dcbxp
```

4 Return to CONFIGURATION mode.

```
exit
```

5 Enable DCBX on all switch ports in CONFIGURATION mode to activate the exchange of DCBX TLV messages with PFC, ETS, and iSCSI configurations.

```
dcbx enable
```

Configure DCBX

View DCBX configuration

```
OS10# show lldp dcbx interface ethernet 1/1/15

E-ETS Configuration TLV enabled                 e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled                r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled                 p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled         f-Application Priority for FCOE disabled
I-Application priority for iSCSI enabled        i-Application Priority for iSCSI disabled
---------------------------------------------------------------------------------------

Interface ethernet1/1/15
    Port Role is Manual
    DCBX Operational Status is Enabled
    Is Configuration Source? FALSE
    Local DCBX Compatibility mode is CEE
    Local DCBX Configured mode is CEE
    Peer Operating version is CEE
    Local DCBX TLVs Transmitted: ErPfI

Local DCBX Status
```

```
-----------------
DCBX Operational Version is 0
DCBX Max Version Supported is 0
Sequence Number: 14
Acknowledgment Number: 5
Protocol State: In-Sync

Peer DCBX Status
-----------------
DCBX Operational Version is 0
DCBX Max Version Supported is 255
Sequence Number: 5
Acknowledgment Number: 14
    220 Input PFC TLV pkts, 350 Output PFC TLV pkts, 0 Error PFC pkts
    220 Input PG TLV Pkts, 396 Output PG TLV Pkts, 0 Error PG TLV Pkts
    71 Input Appln Priority TLV pkts, 80 Output Appln Priority TLV pkts, 0 Error Appln Priority
TLV Pkts

Total DCBX Frames transmitted 538
Total DCBX Frames received 220
Total DCBX Frame errors 0
Total DCBX Frames unrecognized 0
```

View DCBX PFC TLV status

```
OS10# show lldp dcbx interface ethernet 1/1/15 pfc detail

Interface ethernet1/1/15
    Admin mode is on
    Admin is enabled, Priority list is 4,5,6,7
    Remote is enabled, Priority list is 4,5,6,7
    Remote Willing Status is disabled
    Local is enabled, Priority list is 4,5,6,7
    Oper status is init
    PFC DCBX Oper status is Up
    State Machine Type is Feature
    PFC TLV Tx Status is enabled
    Application Priority TLV  Parameters :
    -------------------------------------
    ISCSI TLV Tx Status is enabled
    Local ISCSI PriorityMap is 0x10
    Remote ISCSI PriorityMap is 0x10

    220 Input TLV pkts, 350 Output TLV pkts, 0 Error pkts
    71 Input Appln Priority TLV pkts, 80 Output Appln Priority TLV pkts, 0 Error Appln Priority
TLV Pkts
```

View DCBX ETS TLV status

```
OS10# show lldp dcbx interface ethernet 1/1/15 ets detail

Interface ethernet1/1/15
Max Supported PG is 8
Number of Traffic Classes is 8
Admin mode is on

Admin Parameters :
------------------
Admin is enabled

PG-grp    Priority#         Bandwidth       TSA
-----------------------------------------------
0         0,1,2,3                70%         ETS
1         4,5,6,7                30%         ETS
2                                0%          SP
3                                0%          SP
4                                0%          SP
5                                0%          SP
6                                0%          SP
```

```
7                                   0%          SP
15                                  0%          SP

Remote Parameters :
-------------------
Remote is enabled
PG-grp    Priority#         Bandwidth       TSA
-------------------------------------------------
0         0,1,2,3                 70%         ETS
1         4,5,6,7                 30%         ETS
2                                 0%          SP
3                                 0%          SP
4                                 0%          SP
5                                 0%          SP
6                                 0%          SP
7                                 0%          SP
15                                0%          SP

Remote Willing Status is disabled
Local Parameters :
-------------------
Local is enabled

PG-grp     Priority#         Bandwidth       TSA
-------------------------------------------------
0         0,1,2,3                 70%         ETS
1         4,5,6,7                 30%         ETS
2                                 0%          SP
3                                 0%          SP
4                                 0%          SP
5                                 0%          SP
6                                 0%          SP
7                                 0%          SP
15                                0%          SP

Oper status is init
ETS DCBX Oper status is Up
State Machine Type is Feature
Conf TLV Tx Status is enabled
Reco TLV Tx Status is disabled

220 Input Conf TLV Pkts, 396 Output Conf TLV Pkts, 0 Error Conf TLV Pkts
```

DCBX commands

dcbx enable

Enables DCBX globally on all port interfaces.

Syntax	`dcbx enable`
Parameters	None
Default	Disabled
Command Mode	CONFIGURATION
Usage Information	DCBX is disabled at a global level and enabled at an interface level by default. For DCBX to be operational, DCBX must be enabled at both the global and interface levels. Enable DCBX globally using the `dcbx enable` command to activate the exchange of DCBX TLV messages with PFC, ETS, and iSCSI configurations. To configure the TLVs advertised by a DCBX-enabled port, change the DCBX version, or disable DCBX on an interface, use DCBX interface-level commands. DCBX allows peers to advertise a DCB configuration using LLDP and self-configure with

compatible settings. If you disable DCBX globally on a switch, you can re-enable it to ensure consistent operation of peers in a converged data center network.

Example

```
OS10(config)# dcbx  enable
```

Supported Releases 10.3.0E or later

dcbx tlv-select

Configures the DCB TLVs advertised by a DCBX-enabled port.

Syntax

```
dcbx tlv-select {[ets-conf] [ets-reco] [pfc]}
```

Parameters

- `ets-conf` — Advertise ETS configuration TLVs.
- `ets-reco` — Advertise ETS recommendation TLVs.
- `pfc` — Advertise PFC TLVs.

Default DCBX advertises PFC, ETS Recommendation, and ETS Configuration TLVs.

Command Mode INTERFACE

Usage Information A DCBX-enabled port advertises all TLVs to DCBX peers by default. If PFC or ETS TLVs are disabled, enter the command to re-enable PFC or ETS TLV advertisements. You can enable multiple TLV options, such as ets-conf, ets-reco, and pfc with the same command.

Example

```
OS10(conf-if-eth1/1/2)# dcbx tlv-select ets-conf pfc
```

Supported Releases 10.3.0E or later

dcbx version

Configures the DCBX version used on a port interface.

Syntax

```
dcbx version {auto | cee | ieee}
```

Parameters

- `auto` — Automatically select the DCBX version based on the peer response.
- `cee` — Set the DCBX version to CEE.
- `ieee` — Set the DCBX version to IEEE 802.1Qaz.

Default Auto

Command Mode INTERFACE

Usage Information In Auto mode, a DCBX-enabled port detects an incompatible DCBX version on a peer device port and automatically reconfigures a compatible version on the local port. The `no` version of this command disables the DCBX version.

Example

```
OS10(conf-if-eth1/1/2)# dcbx version cee
```

Supported Releases 10.3.0E or later

lldp tlv-select dcbxp

Enables and disables DCBX on a port interface.

Syntax `lldp tlv-select dcbxp`

Parameters None

Default Enabled interface level; disabled global level

Command Mode INTERFACE

Usage Information DCBX must be enabled at both the global and interface levels. Enable DCBX globally using the `dcbx enable` command to activate the exchange of DCBX TLV messages with PFC, ETS, and iSCSI configurations. To configure the TLVs advertised by a DCBX-enabled port, change the DCBX version, or disable DCBX on an interface, use DCBX interface-level commands . The `no` version of this command disables DCBX on an interface.

Example

```
OS10(conf-if-eth1/1/1)# lldp tlv-select dcbxp
```

Supported Releases 10.3.0E or later

show lldp dcbx interface

Displays the DCBX configuration and PFC or ETS TLV status on an interface.

Syntax `show lldp dcbx interface ethernet` *`node/slot/port[:subport]`* `[ets detail | pfc detail]`

Parameters

- `interface ethernet` *`node/slot/port[:subport]`* — Enter interface information.
- `ets detail` — Display the ETS TLV status and operation with DCBX peers.
- `pfc detail` — Display the PFC TLV status and operation with DCBX peers.

Default Not configured

Command Mode EXEC

Usage Information You must enable DCBX before using this command. DCBX advertises all TLVs — PFC, ETS Recommendation, ETS Configuration, DCBXP, and basic TLVs by default. Enter a port range to display DCBX configuration and TLV operation on multiple ports.

NOTE: In the command output, the `Is configuration source` parameter always displays `False`. `Configuration source` is type of port role that is not supported.

Example (interface)

```
OS10# show lldp dcbx interface ethernet 1/1/15
E-ETS Configuration TLV enabled           e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled          r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled           p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled   f-Application Priority for FCOE disabled
I-Application priority for iSCSI enabled  i-Application Priority for iSCSI disabled
-------------------------------------------------------------------------------

Interface ethernet1/1/15
    Port Role is Manual
    DCBX Operational Status is Enabled
    Is Configuration Source? FALSE
    Local DCBX Compatibility mode is IEEEv2.5
    Local DCBX Configured mode is IEEEv2.5
```

```
    Peer Operating version is IEEEv2.5
    Local DCBX TLVs Transmitted: ERPfI
    5 Input PFC TLV pkts, 2 Output PFC TLV pkts, 0 Error PFC pkts
    5 Input ETS Conf TLV Pkts, 2 Output ETS Conf TLV Pkts, 0 Error ETS Conf TLV
Pkts
    5 Input ETS Reco TLV pkts, 2 Output ETS Reco TLV pkts, 0 Error ETS Reco TLV
Pkts
    5 Input Appln Priority TLV pkts, 2 Output Appln Priority TLV pkts, 0 Error
Appln Priority TLV Pkts

Total DCBX Frames transmitted 8
Total DCBX Frames received 20
Total DCBX Frame errors 0
Total DCBX Frames unrecognized 0
```

Example (ETS detail)

```
OS10# show lldp dcbx interface ethernet 1/1/15 ets  detail
Interface ethernet1/1/15
Max Supported PG is 8
Number of Traffic Classes is 8
Admin mode is on

Admin Parameters :
------------------
Admin is enabled

PG-grp     Priority#         Bandwidth       TSA
-------------------------------------------------
0        0,1,2,3                 70%         ETS
1        4,5,6,7                 30%         ETS
2                                0%          SP
3                                0%          SP
4                                0%          SP
5                                0%          SP
6                                0%          SP
7                                0%          SP

Remote Parameters :
-------------------
Remote is enabled
PG-grp     Priority#         Bandwidth       TSA
-------------------------------------------------
0        0,1,2,3                 70%         ETS
1        4,5,6,7                 30%         ETS
2                                0%          SP
3                                0%          SP
4                                0%          SP
5                                0%          SP
6                                0%          SP
7                                0%          SP

Remote Willing Status is disabled
Local Parameters :
------------------
Local is enabled

PG-grp     Priority#         Bandwidth       TSA
-------------------------------------------------
0        0,1,2,3                 70%         ETS
1        4,5,6,7                 30%         ETS
2                                0%          SP
3                                0%          SP
4                                0%          SP
5                                0%          SP
6                                0%          SP
7                                0%          SP

Oper status is init
ETS DCBX Oper status is Up
```

```
State Machine Type is Asymmetric
Conf TLV Tx Status is enabled
Reco TLV Tx Status is enabled

5 Input Conf TLV Pkts, 2 Output Conf TLV Pkts, 0 Error Conf TLV Pkts
5 Input Reco TLV Pkts, 2 Output Reco TLV Pkts, 0 Error Reco TLV Pkts
```

Example (PFC detail)

```
OS10# show lldp dcbx interface ethernet 1/1/15 pfc detail
Interface ethernet1/1/15
    Admin mode is on
    Admin is enabled, Priority list is 4,5,6,7
    Remote is enabled, Priority list is 4,5,6,7
    Remote Willing Status is disabled
    Local is enabled, Priority list is 4,5,6,7
    Oper status is init
    PFC DCBX Oper status is Up
    State Machine Type is Symmetric
    PFC TLV Tx Status is enabled
    Application Priority TLV  Parameters :
    -------------------------------------
    ISCSI TLV Tx Status is enabled
    Local ISCSI PriorityMap is 0x10
    Remote ISCSI PriorityMap is 0x10

    5 Input TLV pkts, 2 Output TLV pkts, 0 Error pkts
    5 Input Appln Priority TLV pkts, 2 Output Appln Priority TLV pkts, 0 Error
Appln Priority TLV Pkts
```

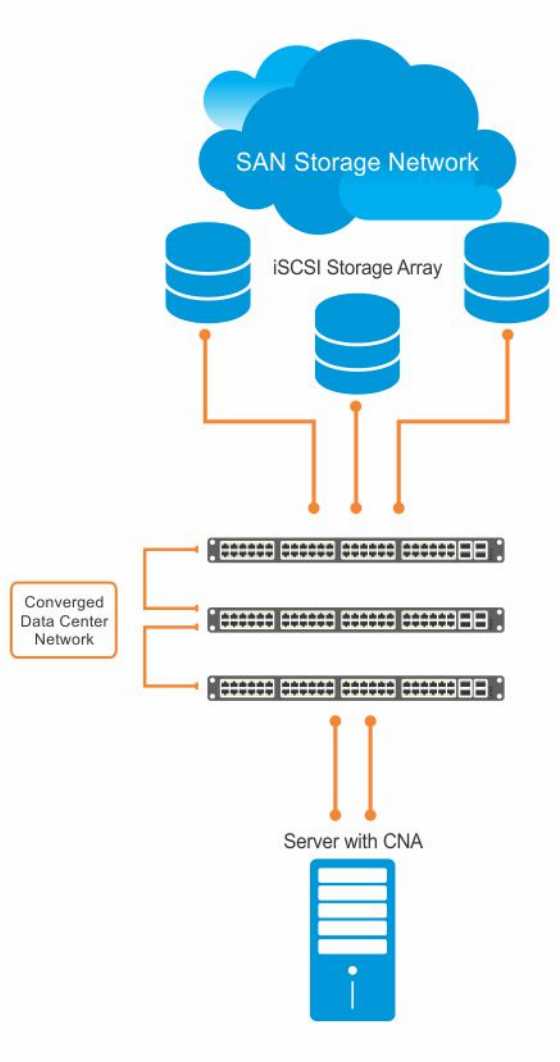
Supported Releases 10.3.0E or later

Internet small computer system interface

iSCSI is a TCP/IP-based protocol that establishes and manages connections between servers and storage devices in a data center network. After you enable iSCSI, iSCSI optimization automatically detects Dell EMC EqualLogic storage arrays directly attached to switch ports. To support storage arrays where auto-detection is not supported, manually configure iSCSI optimization using the `iscsi profile-storage` *`name`* command.

iSCSI optimization enables a switch to auto-detect Dell EMC's iSCSI storage arrays and auto-configure switch ports to improve storage traffic throughput. The switch monitors iSCSI sessions and applies QoS policies on iSCSI traffic. iSCSI optimization operates with or without DCBX over an Ethernet network.

- iSCSI uses the current flow-control configuration by default. If you do not configure flow-control, iSCSI auto-configures flow control settings so that receive-only is enabled and transmit-only is disabled.
- The switch monitors and tracks active iSCSI sessions, including port information and iSCSI session information.
- A user-configured iSCSI CoS profile applies to all iSCSI traffic. Use classifier rules to direct the iSCSI data traffic to queues with preferential QoS treatment over other data passing through the switch. Preferential treatment helps to avoid session interruptions during times of congestion that would otherwise cause dropped iSCSI packets.



In an iSCSI session, a switch connects CNA servers (iSCSI initiators) to a storage array (iSCSI targets) in a SAN or TCP/IP network. iSCSI optimization running on the switch uses dot1p priority-queue assignments to ensure that iSCSI traffic receives priority treatment.

iSCSI configuration notes

- Enable iSCSI optimization so the switch auto-detects and auto-configures Dell EMC EqualLogic storage arrays directly connected to an interface. iSCSI automatically configures switch parameters after connection to a storage device is verified. You must manually enable an interface to support a storage device that is directly connected to a port, but not automatically detected by iSCSI.
- Enable iSCSI session monitoring and the aging time for iSCSI sessions. iSCSI monitoring sessions listen on TCP ports 860 and 3260 by default.
- Configure the CoS/DSCP values applied to ingress iSCSI flows — create a `class-iscsi` class map in POLICY-CLASS-MAP mode.
- Enable LLDP to use iSCSI. The DCBX application TLV carries information about the dot1p priorities to use when sending iSCSI traffic. This informational TLV is packaged in LLDP PDUs. You can reconfigure the 802.1p priority bits advertised in the TLVs.

Configure iSCSI optimization

The iSCSI protocol provides storage traffic TCP/IP transport between servers and storage arrays in a network using iSCSI commands.

1 Configure an interface or interface range to detect a connected storage device.

```
interface ethernet node/slot/port:[subport]

interface range ethernet node/slot/port:[subport]-node/slot/port[:subport]
```

2 Enable the interface to support a storage device that is directly connected to the port and not automatically detected by iSCSI. Use this command for storage devices that do not support LLDP. The switch auto-detects and auto-configures Dell EMC EqualLogic storage arrays directly connected to an interface when you enable iSCSI optimization.

```
iscsi profile-storage storage-device-name
```

3 Configure DCBX to use LLDP to send iSCSI application TLVs with dot1p priorities for iSCSI traffic in INTERFACE mode.

```
lldp tlv-select dcbxp-appln iscsi
```

4 Return to CONFIGURATION mode.

```
exit
```

5 (Optional) If necessary, re-configure the iSCSI TCP ports and IP addresses of target storage devices in CONFIGURATION mode. Separate TCP port numbers with a comma, from 0 to 65535; default 860 and 3260.

```
iscsi target port tcp-port1 [tcp-port2, ..., tcp-port16] [ip-address ip-address]
```

6 Configure the QoS policy applied to ingress iSCSI flows. Apply the service policy to ingress interfaces in CONFIGURATION mode.

(Optional) Reset the default CoS dot1p priority, the default is 4 and/or the trusted DCSP value used for iSCSI traffic. Assign an internal `qos-group` queue, from 0 to 7, to dot1p, from 0 to 7, and DSCP, from 0 to 63, values in POLICY-CLASS-MAP mode.

```
class-map type application class-iscsi
policy-map type application policy-iscsi
   class class-iscsi
      set qos-group traffic-class-number
      set cos dot1p-priority
      set dscp dscp-value
      end
service-policy type application policy-iscsi
```

7 Enable iSCSI monitoring sessions on TCP ports in CONFIGURATION mode.

```
iscsi session-monitoring enable
```

8 (Optional) Set the aging time for the length of iSCSI monitoring sessions in CONFIGURATION mode, 5 to 43,200 minutes; default 10.

```
iscsi aging time [minutes]
```

9 (Optional) Reconfigure the dot1p priority bits advertised in iSCSI application TLVs in CONFIGURATION mode. The default bitmap is `0x10` (dot1p 4). The default dot1p 4 value is sent in iSCSI application TLVs only if you enabled the PFC pause for dot1p 4 traffic using the `pfc-cos dot1p-priority` command.

If you do not configure an `iscsi priority-bits` dot1p value and you configure a `set cos` value in Step 6, the `set cos` value is sent in iSCSI application TLVs. If you configure neither the `iscsi priority-bits` nor the `set cos` value, the default dot1p 4 advertises.

```
iscsi priority-bits dot1p-bitmap
```

10 Enable iSCSI auto-detection and auto-configuration on the switch in CONFIGURATION mode.

```
iscsi enable
```

Configure iSCSI optimization

```
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# iscsi profile-storage compellent
OS10(conf-if-eth1/1/1)# lldp tlv-select dcbxp-appln iscsi
OS10(conf-if-eth1/1/1)# exit

OS10(config)# iscsi target port 3261 ip-address 10.1.1.1
OS10(config)# policy-map type application policy-iscsi
OS10(config-pmap-application)# class class-iscsi
OS10 (config-pmap-c-app)# set qos-group 4
OS10 (config-pmap-c-app)# set cos 4
OS10 (config-pmap-c-app)# exit
OS10(config-pmap-application)# exit

OS10(config)# system qos
OS10(config-sys-qos)# service-policy type application policy-iscsi
OS10(config-sys-qos)# exit
```

```
OS10(config)# iscsi session-monitoring enable
OS10(config)# iscsi aging time 15
OS10(config)# iscsi priority-bits 0x20
OS10(config)# iscsi enable
```

View iSCSI optimization

```
OS10# show iscsi
iSCSI Auto configuration is Enabled
iSCSI session monitoring is  Enabled
iSCSI COS              qos-group 4 remark dot1p 4
Session aging time     15
Maximum number of connections is  100
Port    IP Address
------------------------
3260
860
3261    10.1.1.1
```

```
OS10# show iscsi session detailed
Session 1
---------------------------------------------------
Target:iqn.2001-05.com.equallogic:0-8a0906-00851a00c-98326939fba510a1-517
Initiator:iqn.1991-05.com.microsoft:win-rlkpjo4jun2
Up Time:00:00:18:12(DD:HH:MM:SS)
Time for aging out:29:23:59:35(DD:HH:MM:SS)
ISID:400001370000
Initiator     Initiator  Target       Target    Connection
IP Address    TCP Port   IP Address   TCP Port  ID
-----------------------------------------------------------
10.10.10.210  54748      10.10.10.40  3260      1

Session 2
---------------------------------------------------
Target:iqn.2001-05.com.equallogic:0-8a0906-01251a00c-8ab26939fbd510a1-518
Initiator:iqn.1991-05.com.microsoft:win-rlkpjo4jun2
Up Time:00:00:16:02(DD:HH:MM:SS)
Time for aging out:29:23:59:35(DD:HH:MM:SS)
ISID:400001370000
Initiator     Initiator  Target       Target    Connection
IP Address    TCP Port   IP Address   TCP Port  ID
-----------------------------------------------------------
10.10.10.210  54835      10.10.10.40  3260      1
```

```
OS10# show iscsi storage-devices
Interface Name   Storage Device Name   Auto Detected Status
-------------------------------------------------------------
ethernet1/1/23   EQL-MEM               true
```

iSCSI synchronization on VLT

An iSCSI session is learnt on a VLT LAG during the following scenarios:

- If the iSCSI session receives control packets, as login-request or login-response, on the VLT LAG.
- If the iSCSI session does not receive control packets, but receives data packets on the VLT LAG. This happens when you enable iSCSI session monitoring after the iSCSI session starts.

The information learnt about iSCSI sessions on VLT LAGs synchronizes with the VLT peers.

iSCSI session synchronization happens based on various scenarios:

- If the iSCSI login request is received on an interface that belongs to a VLT LAG, the information synchronizes with the VLT peer and the connection associates with the interface.
- Any updates to connections, including aging updates, that are learnt on VLT LAG members synchronizes with the VLT peer.

- If the iSCSI login request is received on a non-VLT interface, followed by a response from a VLT interface, the connection is associated with the VLT LAG interface and the information about the session synchronizes with the VLT peer.
- When a VLT interconnect comes up, information about iSCSI sessions learnt on the VLT LAG exchanges between the VLT-peers.

iSCSI commands

iscsi aging

Sets the aging time for monitored iSCSI sessions.

Syntax	`iscsi aging [time minutes]`
Parameters	`time minutes` — Enter the aging time in minutes allowed for monitoring iSCSI sessions, from 5 to 43,200.
Default	10 minutes
Command Mode	CONFIGURATION
Usage Information	Configure the aging time allowed for monitored iSCSI sessions on TCP ports before the session closes. The `no` version of this command disables the aging time.
Example	`OS10(config)# iscsi aging time 30`
Supported Releases	10.3.0E or later

iscsi enable

Enables iSCSI auto-detection of attached storage arrays and switch auto-configuration.

Syntax	`iscsi enable`
Parameter	None
Default	None
Command Mode	CONFIGURATION
Usage Information	iSCSI optimization automatically detects storage arrays and auto-configures switch ports with the iSCSI parameters received from a connected device. The `no` version of this command disables iSCSI auto-detection. Starting from release 10.4.1.1, when you perform a fresh installation of OS10, iSCSI autoconfig is enabled and flowcontrol receive is set to on. However, when you upgrade from an earlier release to release 10.4.1.1 or later, the existing iSCSI configuration is retained and the flowcontrol receive could be set to on or off, depending on the iSCSI configuration before upgrade.
Example	`OS10(config)# iscsi enable`
Supported Releases	10.3.0E or later

iscsi priority-bits

Resets the priority bitmap advertised in iSCSI application TLVs.

Syntax	`iscsi priority-bits {priority-bitmap}`
Parameter	*priority-bitmap* — Enter a bitmap value for the dot1p priority advertised for iSCSI traffic in iSCSI application TLVs (0x1 to 0xff).
Default	0x10 (dot1p 4)
Command Mode	CONFIGURATION
Usage Information	iSCSI traffic uses dot1p priority 4 in frame headers by default. Use this command to reconfigure the dot1p-priority bits advertised in iSCSI application TLVs. Enter only one dot1p-bitmap value — setting more than one bitmap value with this command is not supported. The default dot1p 4 value advertises only if you enabled PFC pause frames for dot1p 4 traffic using the `pfc-cos dot1p-priority` command. The `no` version of this command resets to the default value.
Example	`OS10(config)# iscsi priority-bits 0x20`
Supported Releases	10.3.0E or later

iscsi profile-storage

Configures a port for direct connection to a storage device that is not automatically detected by iSCSI.

Syntax	`iscsi profile-storage storage-device-name`
Parameter	*storage-device-name* — Enter a user-defined name of a storage array that iSCSI does not automatically detect.
Default	Not configured
Command Mode	INTERFACE
Usage Information	Configure directly attached storage arrays that iSCSI supports if they are not automatically detected. This command is required for storage devices that do not support LLDP. The `no` version of this command disables the connection.
Example	`OS10(conf-if-eth1/1/2)# iscsi profile-storage compellant`
Supported Releases	10.3.0E or later

iscsi session-monitoring enable

Enables iSCSI session monitoring.

Syntax	`iscsi session-monitoring enable`
Parameter	None
Default	Disabled
Command Mode	CONFIGURATION

Usage Information	To configure the aging timeout in iSCSI monitoring sessions use the `iscsi aging time` command. To configure the TCP ports that listen for connected storage devices in iSCSI monitoring sessions use the `iscsi target port` command. The `no` version of this command disables iSCSI session monitoring. **NOTE: When you enable iSCSI session monitoring, you can monitor a maximum of 100 connections.**
Example	`OS10(config)# iscsi session-monitoring enable`
Supported Releases	10.3.0E or later

iscsi target port

Configures the TCP ports used to monitor iSCSI sessions with target storage devices.

Syntax	`iscsi target port` *`tcp-port1 [tcp-port2, ..., tcp-port16]`* `[ip-address` *`ip-address`*`]`
Parameters	• *`tcp-port`* — Enter one or more TCP port numbers, from 0 to 65535. Separate TCP port numbers with a comma. • `ip-address` *`ip-address`* — (Optional) Enter the IP address in A.B.C.D format of a storage array whose iSCSI traffic is monitored on the TCP port.
Default	3260,860
Command Mode	CONFIGURATION
Usage Information	You can configure a maximum of 16 TCP ports to monitor iSCSI traffic from target storage devices. The `no` version of this command including the IP address removes a TCP port from iSCSI monitoring.
Example	`OS10(config)# iscsi target port 26,40`
Supported Releases	10.3.0E or later

lldp tlv-select dcbxp-appln iscsi

Enables a port to advertise iSCSI application TLVs to DCBX peers.

Syntax	`lldp tlv-select dcbxp-appln iscsi`
Parameter	None
Default	iSCSI application TLVs are advertised to DCBX peers.
Command Mode	INTERFACE
Usage Information	DCB devices use DCBX to exchange iSCSI configuration information with peers and self-configure. iSCSI parameters exchange in time, length, and value (TLV) messages. DCBX requires LLDP enabled to advertise iSCSI application TLVs. iSCSI application TLVs advertise the PFC dot1p priority-bitmap configured using the `iscsi priority-bits` command to DCBX peers. If you do not configure an iSCSI dot1p-bitmap value, iSCSI application TLVs advertise dot1p 4 by default only if you configure dot1p 4 as a PFC priority using the `pfc-cos` command. The `no` version of this command disables iSCSI TLV transmission.
Example	`OS10(conf-if-eth1/1/1)# lldp tlv-select dcbxp-appln iscsi`
Supported Releases	10.3.0E or later

show iscsi

Displays currently configured iSCSI settings.

Syntax `show iscsi`

Parameters None

Command Mode EXEC

Usage Information This command output displays global iSCSI configuration settings. To view target and initiator information use the `show iscsi session` command.

Example

```
OS10# show iscsi
iSCSI Auto configuration is Enabled
iSCSI session monitoring is  Enabled
iSCSI COS              qos-group 4 remark dot1p 4
Session aging time     15
Maximum number of connections is  100
Port    IP Address
------------------------
3260
860
3261    10.1.1.1
```

Supported Releases 10.3.0E or later

show iscsi session

Displays information about active iSCSI sessions.

Syntax `show iscsi session [detailed]`

Parameter `detailed` — Displays a detailed version of the active iSCSI sessions.

Command Mode EXEC

Usage Information In an iSCSI session, `Target` is the storage device, and `Initiator` is the server connected to the storage device.

Example

```
OS10# show iscsi session
```

Example (detailed)

```
OS10# show iscsi session detailed
Session 1
------------------------------------------------
Target:iqn.2001-05.com.equallogic:0-8a0906-00851a00c-98326939fba510a1-517
Initiator:iqn.1991-05.com.microsoft:win-rlkpjo4jun2
Up Time:00:00:18:12(DD:HH:MM:SS)
Time for aging out:29:23:59:35(DD:HH:MM:SS)
ISID:400001370000
Initiator     Initiator  Target       Target    Connection
IP Address    TCP Port   IP Address   TCP Port  ID
------------------------------------------------------------
10.10.10.210  54748      10.10.10.40  3260      1

Session 2
------------------------------------------------
Target:iqn.2001-05.com.equallogic:0-8a0906-01251a00c-8ab26939fbd510a1-518
Initiator:iqn.1991-05.com.microsoft:win-rlkpjo4jun2
Up Time:00:00:16:02(DD:HH:MM:SS)
Time for aging out:29:23:59:35(DD:HH:MM:SS)
ISID:400001370000
Initiator     Initiator  Target       Target    Connection
```

```
IP Address    TCP Port   IP Address   TCP Port  ID
-------------------------------------------------------
10.10.10.210  54835      10.10.10.40  3260      1
```

Supported Releases 10.3.0E or later

show iscsi storage-devices

Displays information about the storage arrays directly attached to OS10 ports.

Syntax `show iscsi storage-devices`

Parameters None

Command Mode EXEC

Usage Information The command output displays the storage device connected to each switch port and whether iSCSI automatically detects it.

Example

```
OS10# show iscsi storage-devices
Interface Name   Storage Device Name   Auto Detected Status
-------------------------------------------------------
ethernet1/1/23   EQL-MEM               true
```

Supported Releases 10.3.0E or later

Converged network DCB example

A converged data center network carries multiple SAN, server, and LAN traffic types that are sensitive to different aspects of data transmission. For example, storage traffic is sensitive to packet loss, while server traffic is latency-sensitive. In a single converged link, all traffic types coexist without imposing serious restrictions on others' performance. DCB allows iSCSI and FCoE SAN traffic to co-exist with server and LAN traffic on the same network. DCB features reduce or avoid dropped frames, retransmission, and network congestion.

DCB provides lossless transmission of FCoE and iSCSI storage traffic using:

- Separate traffic classes for the different service needs of network applications.
- PFC flow control to pause data transmission and avoid dropping packets during congestion.
- ETS bandwidth allocation to guarantee a percentage of shared bandwidth to bursty traffic, while allowing each traffic class to exceed its allocated bandwidth if another traffic class is not using its share.
- DCBX discovery of peers, including PFC, ETS, and other DCB settings parameter exchange, mismatch detection, and remote configuration of DCB parameters.
- iSCSI application protocol TLV information in DCBX advertisements to communicate iSCSI support to peer ports.

This example shows how to configure a DCB converged network in which:

- DCBx is enabled globally to ensure the exchange of DCBx, PFC, ETS, and ISCSI configurations between DCBx-enabled devices.
- PFC is configured to ensure loseless traffic for dot1p priority 4, 5, 6, and 7 traffic.
- ETS allocates 30% bandwidth for dot1p priority 0, 1, 2, and 3 traffic and 70% bandwidth for priority 4, 5, 6, and 7 traffic.
- iSCSI is configured to use dot1p priority 6 for iSCSI traffic, and advertise priority 6 in iSCSI application TLVs.
- The default `class-trust` class map honors dot1p priorities in ingress flows and applies a 1-to-1 dot1p-to-qos-group and a 1-to-1 qos-group-to-queue mapping. In OS10, `qos-group` represents a traffic class used only for internal processing.

1. DCBX configuration (global)

Configure DCBX globally on a switch to enable the exchange of DCBX TLV messages with PFC, ETS, and iSCSI configurations.

```
OS10# configure terminal
OS10(config)# dcbx enable
```

2. PFC configuration (global)

PFC is enabled on traffic classes with dot1p 4, 5, 6, and 7 traffic. All the traffic classes use the default PFC pause settings for shared buffer size and pause frames in ingress queue processing in the network-qos policy map. The pclass policy map honors (trusts) all dot1p ingress traffic. The reserved `class-trust` class map is configured by default. Trust does not modify ingress values in output flows.

```
OS10(config)# class-map type network-qos test4
OS10(config-cmap-nqos)# match qos-group 4
OS10(config-cmap-nqos)# exit
OS10(config)# class-map type network-qos test5
OS10(config-cmap-nqos)# match qos-group 5
OS10(config-cmap-nqos)# exit
OS10(config)# class-map type network-qos test6
OS10(config-cmap-nqos)# match qos-group 6
OS10(config-cmap-nqos)# exit
OS10(config)# class-map type network-qos test7
OS10(config-cmap-nqos)# match qos-group 7
OS10(config-cmap-nqos)# exit

OS10(config)# policy-map type network-qos test
OS10(config-pmap-network-qos)# class test4
OS10(config-pmap-c-nqos)# pause
OS10(config-pmap-c-nqos)# pfc-cos 4
OS10(config-pmap-c-nqos)# exit
OS10(config-pmap-network-qos)# class test5
OS10(config-pmap-c-nqos)# pause
OS10(config-pmap-c-nqos)# pfc-cos 5
OS10(config-pmap-c-nqos)# exit
OS10(config-pmap-network-qos)# class test6
OS10(config-pmap-c-nqos)# pause
OS10(config-pmap-c-nqos)# pfc-cos 6
OS10(config-pmap-c-nqos)# exit
OS10(config-pmap-network-qos)# class test7
OS10(config-pmap-c-nqos)# pause
OS10(config-pmap-c-nqos)# pfc-cos 7
OS10(config-pmap-c-nqos)# exit
OS10(config-pmap-network-qos)# exit

OS10(config)# system qos
OS10(config-sys-qos)# trust-map dscp default
```

3. PFC configuration (interface)

Apply the service policies with dot1p trust and PFC configurations to an interface.

```
OS10(config)# interface ethernet 1/1/53
OS10(conf-if-eth1/1/53)# no shutdown
OS10(conf-if-eth1/1/53)# service-policy input type network-qos test
OS10(conf-if-eth1/1/53)# service-policy input type qos pclass
OS10(conf-if-eth1/1/53)# priority-flow-control mode on
OS10(conf-if-eth1/1/53)# end
```

4. ETS configuration (global)

A trust dot1p-map assigns dot1p 0, 1, 2, and 3 traffic to qos-group 0, and dot1p 4, 5, 6, and 7 traffic to qos-group 1. A `qos-map traffic-` class map assigns the traffic class in qos-group 0 to queue 0, and qos-group 1 traffic to queue 1. A queuing policy map assigns 30% of interface bandwidth to queue 0, and 70% of bandwidth to queue 1.

The pclass policy map applies trust to all dot1p ingress traffic. Trust does not modify ingress dot1p values in output flows. The reserved `class-trust` class map is configured by default.

```
OS10(config)# trust dot1p-map tmap1
OS10(config-tmap-dot1p-map)# qos-group 0 dot1p 0-3
OS10(config-tmap-dot1p-map)# qos-group 1 dot1p 4-7
OS10(config-tmap-dot1p-map)# exit
```

```
OS10(config)# qos-map traffic-class tmap2
OS10(config-qos-map)# queue 0 qos-group 0
OS10(config-qos-map)# queue 1 qos-group 1
OS10(config-qos-map)# exit

OS10(config)# class-map type queuing cmap1
OS10(config-cmap-queuing)# match queue 0
OS10(config-cmap-queuing)# exit
OS10(config)# class-map type queuing cmap2
OS10(config-cmap-queuing)# match queue 1
OS10(config-cmap-queuing)# exit

OS10(config)# policy-map type queuing pmap1
OS10(config-pmap-queuing)# class cmap1
OS10(config-pmap-c-que)# bandwidth percent 30
OS10(config-pmap-c-que)# exit
OS10(config-pmap-queuing)# class cmap2
OS10(config-pmap-c-que)# bandwidth percent 70
OS10(config-pmap-c-que)# end

OS10(config)# system qos
OS10(config-sys-qos)# trust-map dot1p default
```

5. ETS configuration (interface and global)

Apply the service policies with dot1p trust and ETS configurations to an interface or on all switch interfaces. Only one `qos-map traffic-class` map is supported on a switch.

```
OS10(config)# interface ethernet 1/1/53
OS10(conf-if-eth1/1/53)# trust-map dot1p tmap1
OS10(conf-if-eth1/1/53)# qos-map traffic-class tmap2
OS10(conf-if-eth1/1/53)# service-policy input type qos pclass
OS10(conf-if-eth1/1/53)# service-policy output type queuing pmap1
OS10(conf-if-eth1/1/53)# ets mode on
OS10(conf-if-eth1/1/53)# end
```

```
OS10(config)# system qos
OS10(config-sys-qos)# trust-map dot1p tmap1
OS10(config-sys-qos)# qos-map traffic-class tmap2
OS10(config-sys-qos)# service-policy input type qos pclass
OS10(config-sys-qos)# service-policy output type queuing pmap1
OS10(config-sys-qos)# ets mode on
```

6. Verify DCB configuration

```
OS10(conf-if-eth1/1/53)# show configuration
!
interface ethernet1/1/53
 switchport access vlan 1
 no shutdown
 service-policy input type network-qos test
 service-policy input type qos pclass
 service-policy output type queuing pmap1
 ets mode on
 qos-map traffic-class tmap2
 trust-map dot1p tmap1
 priority-flow-control mode on
```

7. Verify DCBX operational status

```
OS10(conf-if-eth1/1/53)# do show lldp dcbx interface ethernet 1/1/53
E-ETS Configuration TLV enabled           e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled          r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled           p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled   f-Application Priority for FCOE disabled
I-Application priority for iSCSI enabled  i-Application Priority for iSCSI disabled
-------------------------------------------------------------------------------

Interface ethernet1/1/53
```

```
 Port Role is Manual
 DCBX Operational Status is Enabled
 Is Configuration Source? FALSE
 Local DCBX Compatibility mode is IEEEv2.5
 Local DCBX Configured mode is AUTO
 Peer Operating version is IEEEv2.5
 Local DCBX TLVs Transmitted: ERPfI
 4 Input PFC TLV pkts, 3 Output PFC TLV pkts, 0 Error PFC pkts
 2 Input ETS Conf TLV Pkts, 27 Output ETS Conf TLV Pkts, 0 Error ETS Conf TLV Pkts
 2 Input ETS Reco TLV pkts, 27 Output ETS Reco TLV pkts, 0 Error ETS Reco TLV Pkts

Total DCBX Frames transmitted 0
Total DCBX Frames received 0
Total DCBX Frame errors 0
Total DCBX Frames unrecognized 0
```

8. Verify PFC configuration and operation

```
OS10(conf-if-eth1/1/53)# do show lldp dcbx interface ethernet 1/1/53 pfc detail

Interface ethernet1/1/53
 Admin mode is on
 Admin is enabled, Priority list is 4,5,6,7
 Remote is enabled, Priority list is 4,5,6,7
 Remote Willing Status is disabled
 Local is enabled, Priority list is 4,5,6,7
 Oper status is init
 PFC DCBX Oper status is Up
 State Machine Type is Symmetric
 PFC TLV Tx Status is enabled
 Application Priority TLV  Parameters :
 -------------------------------------
 ISCSI TLV Tx Status is enabled
 Local ISCSI PriorityMap is 0x10
 Remote ISCSI PriorityMap is 0x10

 4 Input TLV pkts, 3 Output TLV pkts, 0 Error pkts
 4 Input Appln Priority TLV pkts, 3 Output Appln Priority TLV pkts,
 0 Error Appln Priority TLV Pkts
```

9. Verify ETS configuration and operation

```
OS10(conf-if-eth1/1/53)# do show lldp dcbx interface ethernet 1/1/53 ets detail

Interface ethernet1/1/53
Max Supported PG is 8
Number of Traffic Classes is 8
Admin mode is on

Admin Parameters :
------------------
Admin is enabled

PG-grp     Priority#         Bandwidth       TSA
-----------------------------------------------
0        0,1,2,3,                30%         ETS
1        4,5,6,7                 70%         ETS
2                                0%          ETS
3                                0%          ETS
4                                0%          ETS
5                                0%          ETS
6                                0%          ETS
7                                0%          ETS

Remote Parameters :
-------------------
Remote is enabled
PG-grp     Priority#         Bandwidth       TSA
-----------------------------------------------
```

```
0        0,1,2,3,                  30%          ETS
1        4,5,6,7                   70%          ETS
2                                  0%           SP
3                                  0%           SP
4                                  0%           SP
5                                  0%           SP
6                                  0%           SP
7                                  0%           SP

Remote Willing Status is disabled
Local Parameters :
-------------------
Local is enabled

PG-grp     Priority#         Bandwidth       TSA
-----------------------------------------------
0        0,1,2,3,                  30%          ETS
1        4,5,6,7                   70%          ETS
2                                  0%           ETS
3                                  0%           ETS
4                                  0%           ETS
5                                  0%           ETS
6                                  0%           ETS
7                                  0%           ETS

Oper status is init
ETS DCBX Oper status is Up
State Machine Type is Asymmetric
Conf TLV Tx Status is enabled
Reco TLV Tx Status is enabled

2 Input Conf TLV Pkts, 27 Output Conf TLV Pkts, 0 Error Conf TLV Pkts
2 Input Reco TLV Pkts, 27 Output Reco TLV Pkts, 0 Error Reco TLV Pkts
```

10. iSCSI optimization configuration (global)

This example accepts the default settings for aging time and TCP ports used in monitored iSCSi sessions. A Compellant storage array is connected to the port. The policy-iscsi policy map sets the CoS dot1p priority used for iSCSI traffic to `6` globally on the switch. By default, iSCSI traffic uses priority 4. The `iscsi priority-bits 0x40` command sets the advertised dot1p priority used by iSCSI traffic in application TLVs to `6`. Hexadecimal `0x40` is binary `0 1 0 0 0 0 0 0`.

```
OS10(conf-if-eth1/1/53)# iscsi profile-storage compellent
OS10(conf-if-eth1/1/53)# lldp tlv-select dcbxp-appln iscsi
OS10(conf-if-eth1/1/53)# exit

OS10(config)# iscsi target port 3261 ip-address 10.1.1.1
OS10(config)# policy-map type application policy-iscsi
OS10(config-pmap-application)# class class-iscsi
OS10(config-pmap-c-app)# set qos-group 6
OS10(config-pmap-c-app)# set cos 6
OS10(config-pmap-c-app)# exit
OS10(config-pmap-application)# exit

OS10(config)# system qos
OS10(config-sys-qos)# service-policy type application policy-iscsi
OS10(config-sys-qos)# exit

OS10(config)# iscsi session-monitoring enable
OS10(config)# iscsi priority-bits 0x40
OS10(config)# iscsi enable
```

11. Verify iSCSI optimization (global)

After you enable iSCSI optimization, the iSCSI application priority TLV parameters are added in the `show` command output to verify a PFC configuration.

```
OS10(conf-if-eth1/1/53)# do show lldp dcbx interface ethernet 1/1/53 pfc detail

Interface ethernet1/1/53
    Admin mode is on
    Admin is enabled, Priority list is 4,5,6,7
    Remote is enabled, Priority list is 4,5,6,7
    Remote Willing Status is disabled
    Local is enabled, Priority list is 4,5,6,7
    Oper status is init
    PFC DCBX Oper status is Up
    State Machine Type is Symmetric
    PFC TLV Tx Status is enabled
    Application Priority TLV  Parameters :
    -------------------------------------
    ISCSI TLV Tx Status is enabled
    Local ISCSI PriorityMap is 0x40
    Remote ISCSI PriorityMap is 0x10

    4 Input TLV pkts, 3 Output TLV pkts, 0 Error pkts
    4 Input Appln Priority TLV pkts, 3 Output Appln Priority TLV pkts, 0 Error Appln Priority
TLV Pkts
```

12. DCBX configuration (interface)

This example shows how to configure and verify different DCBX versions.

```
OS10(conf-if-eth1/1/53)# dcbx version cee
OS10(conf-if-eth1/1/53)# show configuration
!
interface ethernet1/1/53
 switchport access vlan 1
 no shutdown
 dcbx version cee
 service-policy input type network-qos test
 service-policy input type qos pclass
 service-policy output type queuing pmap1
 ets mode on
 qos-map traffic-class tmap2
 trust-map dot1p tmap1
 priority-flow-control mode on

OS10(conf-if-eth1/1/53)# do show lldp dcbx interface ethernet 1/1/53
E-ETS Configuration TLV enabled           e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled          r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled           p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled   f-Application Priority for FCOE disabled
I-Application priority for iSCSI enabled  i-Application Priority for iSCSI disabled
---------------------------------------------------------------------------------

Interface ethernet1/1/53
 Port Role is Manual
 DCBX Operational Status is Enabled
 Is Configuration Source? FALSE
 Local DCBX Compatibility mode is CEE
 Local DCBX Configured mode is CEE
 Peer Operating version is CEE
 Local DCBX TLVs Transmitted: ErPfi

Local DCBX Status
-----------------
DCBX Operational Version is 0
DCBX Max Version Supported is 0
Sequence Number: 2
Acknowledgment Number: 1
Protocol State: In-Sync
```

```
Peer DCBX Status
-----------------
DCBX Operational Version is 0
DCBX Max Version Supported is 0
Sequence Number: 1
Acknowledgment Number: 2
 3 Input PFC TLV pkts, 3 Output PFC TLV pkts, 0 Error PFC pkts
 3 Input PG TLV Pkts, 3 Output PG TLV Pkts, 0 Error PG TLV Pkts
 3 Input Appln Priority TLV pkts, 3 Output Appln Priority TLV pkts,
 0 Error Appln Priority TLV Pkts

Total DCBX Frames transmitted 3
Total DCBX Frames received 3
Total DCBX Frame errors 0
Total DCBX Frames unrecognized
0
```

```
OS10(conf-if-eth1/1/53)# dcbx version cee
OS10(conf-if-eth1/1/53)# show configuration
!
interface ethernet1/1/53
 switchport access vlan 1
 no shutdown
 dcbx version ieee
 service-policy input type network-qos test
 service-policy input type qos pclass
 service-policy output type queuing pmap1
 ets mode on
 qos-map traffic-class tmap2
 trust-map dot1p tmap1
 priority-flow-control mode on

OS10(conf-if-eth1/1/53)# do show lldp dcbx interface ethernet 1/1/53
E-ETS Configuration TLV enabled           e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled          r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled           p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled   f-Application Priority for FCOE disabled
I-Application priority for iSCSI enabled  i-Application Priority for iSCSI disabled
---------------------------------------------------------------------------------

Interface ethernet1/1/53
 Port Role is Manual
 DCBX Operational Status is Enabled
 Is Configuration Source? FALSE
 Local DCBX Compatibility mode is IEEEv2.5
 Local DCBX Configured mode is IEEEv2.5
 Peer Operating version is IEEEv2.5
 Local DCBX TLVs Transmitted: ERPfI
 13 Input PFC TLV pkts, 4 Output PFC TLV pkts, 0 Error PFC pkts
 3 Input ETS Conf TLV Pkts, 26 Output ETS Conf TLV Pkts, 0 Error ETS Conf TLV Pkts
 3 Input ETS Reco TLV pkts, 26 Output ETS Reco TLV pkts, 0 Error ETS Reco TLV Pkts

Total DCBX Frames transmitted 0
Total DCBX Frames received 0
Total DCBX Frame errors 0
Total DCBX Frames unrecognized 0
```

16

sFlow

sFlow is a standard-based sampling technology embedded within switches and routers that monitors network traffic. It provides traffic monitoring for high-speed networks with many switches and routers.

- OS10 supports sFlow version 5
- Only data ports support sFlow collector
- OS10 supports a maximum of two sFlow collectors
- OS10 does not support sFlow on SNMP, VLAN, tunnel interfaces, extended sFlow, backoff mechanism, and egress sampling

sFlow uses two types of sampling:

- Statistical packet-based sampling of switched or routed packet flows
- Time-based sampling of interface counters

NOTE: On the S4248FB-ON and the S4248FBL-ON platforms, sampling is performed based on the cumulative packet counts from all the sFlow enabled ports.

sFlow monitoring consists of an sFlow agent embedded in the device and an sFlow collector:

- The sFlow agent resides anywhere within the path of the packet. The agent combines the flow samples and interface counters into sFlow datagrams and forwards them to the sFlow collector at regular intervals. The datagrams consist of information on, but not limited to, the packet header, ingress and egress interfaces, sampling parameters, and interface counters. Application-specific integrated circuits (ASICs) handle packet sampling.
- The sFlow collector analyses the datagrams received from different devices and produces a network-wide view of traffic flows.

Enable sFlow

You can enable sFlow either on all interfaces globally or on a specific set of interfaces. The system displays an error message if you try to enable sFlow on both modes at one time.

If you configure sFlow only on a set of interfaces, any further change to the sFlow-enabled ports triggers the sFlow agent to restart. This results in a gap in the polling counter statistics of 30 seconds and the sFlow counters are reset on all sFlow-enabled ports.

When you enable sFlow on a port-channel:

- When in Per-Interface mode, the counter statistics of sFlow-enabled ports reset to zero when you add a new member port or remove an existing member port from any sflow enabled port-channel group.
- sFlow counter statistics that are individually reported for the port members of a port-channel data source are accurate. Counter statistics reported for the port-channel may not be accurate. To calculate the correct counters for a port-channel data source, add together the counter statistics of the individual port members.

Enable or disable sFlow globally

sFlow is disabled globally by default.

- Enable sFlow globally on all interfaces in CONFIGURATION mode.
  ```
  sflow enable all-interfaces
  ```
- Disable sFlow in CONFIGURATION mode.
  ```
  no sflow
  ```

Enable or disable sFlow on a specific interface

- Enable sFlow in CONFIGURATION mode.

  ```
  sflow enable
  ```

- Disable sFlow in CONFIGURATION mode.

  ```
  no sflow enable
  ```

Enable sFlow on a specific interface

```
OS10(config)# sflow enable
OS10(config)# interface ethernet 1/1/1
OS10(conf-if-eth1/1/1)# sflow enable
```

Enable sFlow on a range of interfaces

```
OS10(config)# sflow enable
OS10(config)# interface range ethernet 1/1/1-1/1/10
OS10(conf-range-eth1/1/1-1/1/10)# sflow enable
```

Enable sFlow on a port-channel

```
OS10(config)# sflow enable
OS10(config)# interface range port-channel 1-10
OS10(conf-range-po-1-10)# sflow enable
```

Max-header size configuration

- Set the packet maximum size in CONFIGURATION mode, from 64 to 256. The default is 128 bytes.

  ```
  max-header-size header-size
  ```

- Disable the header size in CONFIGURATION mode.

  ```
  no sflow max-header-size
  ```

- View the maximum packet header size in EXEC mode.

  ```
  show sflow
  ```

Configure sFlow maximum header size

```
OS10(config)# sflow max-header-size 80
```

View sFlow information

```
OS10# show sflow
sFlow services are enabled
Management Interface sFlow services are disabled
Global default sampling rate: 32768
Global default counter polling interval: 20
Global default extended maximum header size: 128 bytes
Global extended information enabled: none
1 collector(s) configured
Collector IP addr:10.16.151.245 Agent IP addr:10.16.132.181 UDP port:6343 VRF:Default
31722 UDP packets exported
0 UDP packets dropped
34026 sFlow samples collected
```

View sFlow running configuration

```
OS10# show running-configuration sflow
sflow enable
sflow max-header-size 80
sflow polling-interval 30
sflow sample-rate 4096
sflow collector 10.16.150.1 agent-addr 10.16.132.67 6767 max-datagram-size 800
sflow collector 10.16.153.176 agent-addr 3.3.3.3 6666
!
interface ethernet1/1/1
```

```
sflow enable
!
```

Collector configuration

Configure the IPv4 or IPv6 address for the sFlow collector. When you configure the collector, enter a valid and reachable IPv4 or IPv6 address. You can configure a maximum of two sFlow collectors. If you specify two collectors, samples are sent to both. The agent IP address must be the same for both the collectors.

- Enter an IPv4 or IPv6 address for the sFlow collector, IPv4 or IPv6 address for the agent, UDP collector port number, maximum datagram size, and the VRF instance in CONFIGURATION mode.

  ```
  sflow collector {ip-address | ipv6-address} agent-addr {ip-address | ipv6-address} [collector-port-number] [vrf vrf-name]
  ```

The `no` form of the command disables sFlow collectors in CONFIGURATION mode.

If you configure a collector for a nondefault VRF, create the VRF first. If you do not specify the VRF instance, the system configures the collector for the default VRF instance.

sFlow collector

```
OS10(config)# sflow collector 10.1.1.1 agent-addr 2.2.2.2 6343  vrf red
```

View sFlow information

```
OS10# show sflow
sFlow services are enabled
Management Interface sFlow services are disabled
Global default sampling rate: 32768
Global default counter polling interval: 10
Global default extended maximum header size: 128 bytes
Global extended information enabled: none
2 collector(s) configured
Collector IP addr:10.1.1.1 Agent IP addr:2.2.2.2 UDP port:6343 VRF:red
0 UDP packets exported
0 UDP packets dropped
0 sFlow samples collected
```

Polling-interval configuration

The polling interval for an interface is the number of seconds between successive samples of counters sent to the collector. You can configure the duration for polled interface statistics. Unless there is a specific deployment need to configure a lower polling interval value, configure the polling interval to the maximum value.

- Change the default counter polling interval in CONFIGURATION mode, from 10 to 300. The default is 20.

  ```
  sflow polling-interval interval-size
  ```

- Disable the polling interval in CONFIGURATION mode.

  ```
  no sflow polling-interval
  ```

- View the polling interval in EXEC mode.

  ```
  show sflow
  ```

Configure sFlow polling interval

```
OS10(config)# sflow polling-interval 200
```

View sFlow information

```
OS10# show sflow
sFlow services are enabled
Management Interface sFlow services are disabled
Global default sampling rate: 32768
```

```
Global default counter polling interval: 200
Global default extended maximum header size: 128 bytes
Global extended information enabled: none
1 collector(s) configured
Collector IP addr:10.16.151.245 Agent IP addr:10.16.132.181 UDP port:6343 VRF:Default
31722 UDP packets exported
0 UDP packets dropped
34026 sFlow samples collected
```

View sFlow running configuration

```
OS10# show running-configuration sflow
sflow enable
sflow max-header-size 80
sflow polling-interval 200
sflow sample-rate 4096
sflow collector 10.16.150.1 agent-addr 10.16.132.67 6767 max-datagram-size 800
sflow collector 10.16.153.176 agent-addr 3.3.3.3 6666
!
interface ethernet1/1/1
sflow enable
!
```

Sample-rate configuration

Sampling rate is the number of packets skipped before the sample is taken. If the sampling rate is 4096, one sample generates for every 4096 packets observed.

- Set the sampling rate in CONFIGURATION mode, from 4096 to 65535. The default is 32768.

  ```
  sflow sample-rate sampling-size
  ```

- Disable packet sampling in CONFIGURATION mode.

  ```
  no sflow sample-rate
  ```

- View the sampling rate in EXEC mode.

  ```
  show sflow
  ```

Configure sFlow sampling rate

```
OS10(config)# sflow sample-rate 4096
```

View sFlow packet header size

```
OS10# show sflow
sFlow services are enabled
Management Interface sFlow services are disabled
Global default sampling rate: 4096
Global default counter polling interval: 20
Global default extended maximum header size: 128 bytes
Global extended information enabled: none
1 collector(s) configured
Collector IP addr:10.16.151.245 Agent IP addr:10.16.132.181 UDP port:6343 VRF:Default
31722 UDP packets exported
0 UDP packets dropped
34026 sFlow samples collected
```

View sFlow running configuration

```
OS10# show running-configuration sflow
sflow enable
sflow max-header-size 80
sflow polling-interval 20
sflow sample-rate 4096
sflow collector 10.16.150.1 agent-addr 10.16.132.67 6767 max-datagram-size 800
sflow collector 10.16.153.176 agent-addr 3.3.3.3 6666
!
interface ethernet1/1/1
```

```
sflow enable
!
```

Source interface configuration

You can configure an interface as a source for sFlow. The sFlow agent uses the IP address of the configured source interface as the agent IP address.

- Configure the source interface in CONFIGURATION mode.

  ```
  sflow source-interface {ethernet node/slot/port[:subport] | loopback loopback-ID| port-channel port-channel-ID| vlan vlan-ID}
  ```

- View the interface details.

  ```
  show running-configuration sflow
  ```

  ```
  show sflow
  ```

Configure sFlow source interface

```
OS10(config)# sflow source-interface ethernet 1/1/1

OS10(config)# sflow source-interface port-channel 1

OS10(config)# sflow source-interface loopback 1

OS10(config)# sflow source-interface vlan 10
```

View sFlow running configuration

```
OS10# show running-configuration sflow
sflow enable all-interfaces
sflow source-interface vlan10
sflow collector 5.1.1.1 agent-addr 4.1.1.1 6343
sflow collector 6.1.1.1 agent-addr 4.1.1.1 6343

OS10(config)#show running-configuration interface vlan
!
interface vlan1
 no shutdown
!
interface vlan10
 no shutdown
 ip address 10.1.1.1/24
```

View sFlow details

```
OS10# show sflow
sFlow services are enabled
Management Interface sFlow services are disabled
Global default sampling rate: 32768
Global default counter polling interval: 30
Global default extended maximum header size: 128 bytes
Global extended information enabled: none
2 collector(s) configured
Collector IP addr:5.1.1.1 Agent IP addr:10.1.1.1 UDP port:6343 VRF:Default → It shows active
agent-ip
Collector IP addr:6.1.1.1 Agent IP addr:10.1.1.1 UDP port:6343 VRF:Default → It shows active
agent-ip
2 UDP packets exported
0 UDP packets dropped
2 sFlow samples collected
```

View sFlow information

OS10 does not support statistics for UDP packets dropped and samples received from the hardware.

- View sFlow configuration details and statistics in EXEC mode.

```
OS10# show sflow
sFlow services are enabled
Management Interface sFlow services are disabled
Global default sampling rate: 32768
Global default counter polling interval: 30
Global default extended maximum header size: 128 bytes
Global extended information enabled: none
1 collector(s) configured
Collector IP addr:10.16.151.245 Agent IP addr:10.16.132.181 UDP port:6343 VRF:Default
31722 UDP packets exported
0 UDP packets dropped
34026 sFlow samples collected
```

- View sFlow configuration details on a specific interface in EXEC mode.

```
OS10# show sflow interface port-channel 1
port-channel1
sFlow is enabled on port-channel1
Samples rcvd from h/w: 0
```

- View the sFlow running configuration in EXEC mode.

```
OS10# show running-configuration sflow
sflow enable
sflow max-header-size 80
sflow polling-interval 30
sflow sample-rate 4096
sflow collector 10.16.150.1 agent-addr 10.16.132.67 6767
sflow collector 10.16.153.176 agent-addr 3.3.3.3 6666
!
interface ethernet1/1/1
sflow enable
!
```

sFlow commands

sflow collector

Configures an sFlow collector IP address where sFlow datagrams are forwarded. You can configure a maximum of two collectors.

Syntax `sflow collector {ipv4-address | ipv6-address} agent-addr {ipv4-address | ipv6-address} [collector-port-number] [vrf vrf-name]`

Parameters

- *ipv4-address | ipv6-address* — Enter an IPv4 or IPv6 address in A.B.C.D/A::B format.
- `agent-addr` *ipv4-address | ipv6-address* — Enter the sFlow agent IP address. If you configure two collectors, the agent IP address must be the same for both the collectors.
- *collector-port-number* — (Optional) Enter the UDP port number, from 1 to 65535. The default is 6343.
- `vrf` — (Optional) Enter the VRF instance to set the VRF context to the collector IP address. If you do not specify a VRF, the system uses the default VRF.

Defaults Not configured

Command Modes CONFIGURATION

Usage Information	You must enter a valid and reachable IPv4 or IPv6 address. If you configure two collectors, traffic samples are sent to both. The sFlow agent address is the IPv4 or IPv6 address used to identify the agent to the collector. The `no` version of this command removes the configured sFlow collector. If you specify a nondefault VRF with the `sflow collector` command, configure the VRF first.
Example	`OS10(conf)# sflow collector 10.1.1.1 agent-addr 2.2.2.2 6343vrf default` `OS10(conf)# sflow collector 10.1.1.1 agent-addr 2.2.2.2 6343vrf vrf-core`
Supported Releases	10.3.0E or later. Updated the command to specify a nondefault VRF on OS10 release 10.4.3.0 or later

sflow enable

Enables sFlow on a specific interface or globally on all interfaces.

Syntax	`sflow enable [all-interfaces]`
Parameters	`all-interfaces` — (Optional) Enter to enable sFlow globally.
Default	Disabled
Command Mode	CONFIGURATION
Usage Information	The `no` version of this command to disables sFlow.

Example (interface)	`OS10(config)# sflow enable`
	`OS10(config)# interface ethernet 1/1/1`
	`OS10(conf-if-eth1/1/1)# sflow enable`

Example (interface range)	`OS10(config)# sflow enable`
	`OS10(config)# interface range ethernet 1/1/1-1/1/10`
	`OS10(conf-range-eth1/1/1-1/1/10)# sflow enable`

Example (port-channel)	`OS10(config)# sflow enable`
	`OS10(config)# interface range port-channel 1-10`
	`OS10(conf-range-po-1-10)# sflow enable`

Supported Releases	10.3.0E or later

sflow max-header-size

Sets the maximum header size of a packet.

Syntax	`sflow max-header-size` *`header-size`*
Parameter	*`header-size`* — Enter the header size in bytes, from 64 to 256. The default is 128.
Default	128 bytes
Command Mode	CONFIGURATION
Usage Information	Use the `no` version of the command to reset the header size to the default value.

Example	`OS10(conf)# sflow max-header-size 256`
Supported Releases	10.3.0E or later

sflow polling-interval

Sets the sFlow polling interval.

Syntax	`sflow polling-interval` *`interval-value`*
Parameter	*`interval-value`* — Enter the interval value in sections, from 10 to 300. The default is 30.
Defaults	30
Command Mode	CONFIGURATION
Usage Information	The polling interval for an interface is the number of seconds between successive samples of counters sent to the collector. You can configure the duration for polled interface statistics. The `no` version of the command resets the interval time to the default value.
Example	`OS10(conf)# sflow polling-interval 200`
Supported Releases	10.3.0E or later

sflow sample-rate

Configures the sampling rate.

Syntax	`sflow sample-rate` *`value`*
Parameter	*`value`* — Enter the packet sample rate, from 4096 to 65535. The default is 32768.
Default	32768
Command Mode	CONFIGURATION
Usage Information	Sampling rate is the number of packets skipped before the sample is taken. For example, if the sampling rate is 4096, one sample generates for every 4096 packets observed. The `no` version of the command resets the sampling rate to the default value.
Example	`OS10(conf)# sflow sample-rate 4096`
Supported Releases	10.3.0E or later

sflow source-interface

Configures an interface as source for sFlow. The sFlow agent uses the IP address of the configured source interface as the agent IP address.

Syntax	`sflow source-interface {ethernet` *`node/slot/port[:subport]`* `| loopback` *`loopback-ID`* `| port-channel` *`port-channel-ID`* `| vlan` *`vlan-ID`*`}`

Parameters

- `ethernet node/slot/port[:subport]`—Enter the physical interface type details.
- `loopback loopback-ID`—Enter the Loopback interface details. The Loopback ID range is from 0 to 16383.
- `port-channel port-channel-ID`—Enter the port channel details. The port channel ID range is from 1 to 128.
- `vlan vlan-ID`—Enter the VLAN details. The VLAN ID range is from 1 to 4093.

Default Disabled

Command Mode CONFIGURATION

Usage Information The `no` version of this command removes the configuration from the interface.

Example (Ethernet)
```
OS10(config)# sflow source-interface ethernet 1/1/1
```

Example (Loopback)
```
OS10(config)# sflow source-interface loopback 1
```

Example (port-channel)
```
OS10(config)# sflow source-interface port-channel 1
```

Example (VLAN)
```
OS10(config)# sflow source-interface vlan 10
```

Supported Releases 10.4.1.0 or later

show sflow

Displays the current sFlow configuration for all interfaces or by a specific interface type.

Syntax `show sflow [interface type]`

Parameter `interface type` — (Optional) Enter either `ethernet` or `port-channel` for the interface type.

Command Mode EXEC

Usage Information OS10 does not support statistics for UDP packets dropped and samples received from the hardware.

Example
```
OS10# show sflow
sFlow services are enabled
Management Interface sFlow services are disabled
Global default sampling rate: 32768
Global default counter polling interval: 30
Global default extended maximum header size: 128 bytes
Global extended information enabled: none
1 collector(s) configured
Collector IP addr:10.16.151.245 Agent IP addr:10.16.132.181 UDP port:6343
VRF:Default
31722 UDP packets exported
0 UDP packets dropped
34026 sFlow samples collected
```

```
OS10# show sflow
sFlow services are enabled
Management Interface sFlow services are disabled
Global default sampling rate: 32768
Global default counter polling interval: 30
Global default extended maximum header size: 128 bytes
Global extended information enabled: none
1 collector(s) configured
Collector IP addr:10.16.151.145 Agent IP addr:10.16.132.160 UDP port:6343
```

```
VRF:RED
0 UDP packets exported
0 UDP packets dropped
0 sFlow samples collected
```

Example (port-channel)

```
OS10# show sflow interface port-channel 1
port-channel1
sFlow is enabled on port-channel1
Samples rcvd from h/w: 0
```

Supported Releases 10.3.0E or later. Updated the command to show the nondefault VRF

17

Telemetry

Network health relies on performance monitoring and data collection for analysis and troubleshooting. Network data is often collected with SNMP and CLI commands using the pull mode. In pull mode, a management device sends a get request and pulls data from a client. As the number of objects in the network and the metrics grow, traditional methods limit network scaling and efficiency. Using multiple management systems further limits network scaling. The pull model increases the processing load on a switch by collecting all data even when there is no change.

Streaming telemetry provides an alternative method where data is continuously transmitted from network devices with efficient, incremental updates. Operators subscribe to the specific data they need using well-defined sensor identifiers.

While SNMP management systems poll for data even if there is no change, streaming telemetry enables access to near real-time, model-driven, and analytics-ready data. It supports more effective network automation, traffic optimization, and preventative troubleshooting.

For example, streaming telemetry reports packet drops or high utilization on links in real time. A network automation application can use this information to provision new paths and optimize traffic transmission across the network. The data is encoded using Google Protocol Buffers (GPB) and streamed using Google Protocol RPC (gRPC) transport.

You can use OS10 telemetry to stream data to:

- Dell-implemented external collectors, such as VMware vRNI or Wavefront
- Proprietary network collectors that you implement

Telemetry terminology

Dial-out mode	The switch initiates a session with one or more devices according to the sensor paths and destinations in a subscription.
Sensor path	The path used to collect data for streaming telemetry.
Sensor group	A reusable group of multiple sensor paths and exclude filters.
Destination group	The IP address and transport port on a destination server to which telemetry data is streamed. You can configure multiple destinations and reuse the destination group in subscription profiles.
Subscription profile	Data collector destinations and stream attributes that are associated with sensor paths. A subscription ties sensor paths and a destination group with a transport protocol, encoding format, and streaming interval. The telemetry agent in the switch attempts to establish a session with each collector in the subscription profile, and streams data to the collector. If a collector is not reachable, the telemetry agent continuously tries to establish the connection at one-minute intervals.

YANG-modeled telemetry data

This section describes the YANG containers from which telemetry data can be streamed to destinations with the recommended minimum sampling intervals.

BGP

Table 71. BGP

YANG Container	Minimum sampling interval (milliseconds)
bgp/bgp-oper/bgpPeerCount	15000
bgp/bgp-oper/bgpPrfxCntrsEntry	15000

BGP peers

Table 72. BGP peers

YANG Container	Minimum sampling interval (milliseconds)
infra-bgp/peer-state/peer-status	0

Buffer statistics

Table 73. Buffer statistics

YANG Container	Minimum sampling interval (milliseconds)
base-qos/queue-stat	15000
base-qos/priority-group-stat	15000
base-qos/buffer-pool-stat	15000
base-qos/buffer-pool	15000

Device information

Table 74. Device information

YANG Container	Minimum sampling interval (milliseconds)
base-pas/chassis	15000
base-pas/card	15000
base-switch/switching-entities/switch-stats	15000

Environmental statistics

Table 75. Environmental statistics

YANG Container	Minimum sampling interval (milliseconds)
base-pas/entity	15000
base-pas/psu	15000
base-pas/fan-try	15000
base-pas/fan	15000
base-pas/led	15000
base-pas/temperature	15000
base-pas/temp_threshold	15000
base-pas/media	15000
base-pas/media-channel	15000

Interface statistics

Table 76. Interface statistics

YANG Container	Minimum sampling interval (milliseconds)
if/interfaces-state/interface/statistics	15000
dell-base-if-cmn/if/interfaces-state/interface	15000

Port-channel (lag) member ports

Table 77. Port-channel (lag) member ports

YANG Container	Minimum sampling interval (milliseconds)
dell-base-if-cmn/if/interfaces	0

System statistics

Table 78. System statistics

YANG Container	Minimum sampling interval (milliseconds)
system-status/current-status	15000

Configure telemetry

To enable the streaming of telemetry data to destinations in a subscription profile:

1 Enable telemetry on the switch.
2 Configure a destination group.
3 Configure a subscription profile by associating one or more destination groups and pre-configured sensor groups.

After you complete Step 3, the telemetry agent starts streaming data to destination devices.

Configuration notes

- The telemetry agent collects data from OS10 applications and switch hardware. When you configure a sampling rate of 0, which is near real-time, telemetry collects data as soon as an event occurs. If you configure a sampling rate, telemetry performs periodic data collection. The recommended minimum sampling intervals are described in **Configure a sensor group**.
- OS10 telemetry supports:
  - Only one configured destination group, and only one destination address in the group.
  - Only one subscription profile.

Enable telemetry

1 Enter telemetry mode from CONFIGURATION mode.
```
OS10(config)# telemetry
```
2 Enable streaming telemetry in TELEMETRY mode.
```
OS10(conf-telemetry)# enable
```

Configure a sensor group

A sensor group defines the data that is collected and streamed to a destination. Use any of the pre-configured sensor groups to monitor system resources. To display the sensor paths for each group, use the `show telemetry sensor-group` command.

Table 79. Pre-configured sensor group

Pre-configured sensor group	Minimum sampling interval (milliseconds)
BGP	15000
BGP-peer	0
Buffer	15000
Device	15000
Environment	15000
Interface	15000
LAG (port channel)	0
System	15000

Configure a destination group

A destination group defines the destination servers to which streaming telemetry data is sent.

1 Enter the destination group name in TELEMETRY mode. A maximum of 32 characters.

```
OS10(conf-telemetry)# destination-group group-name
```

2 Enter the IPv4 or IPv6 address and transport-service port number in DESTINATION-GROUP mode. Only one destination is supported in the 10.4.3.0 release. You can enter a fully qualified domain name (FQDN) for *ip-address*. The destination domain name resolves to an IP address — see System domain name and list.

```
OS10(conf-telemetry-dg-dest)# destination ip-address port-number
```

3 Return to TELEMETRY mode.

```
OS10(conf-telemetry-dg-dest)# exit
```

Configure a subscription profile

A subscription profile associates destination groups and sensor groups, and specifies the data encoding format and transport protocol.

1 Enter the subscription profile name in TELEMETRY mode. A maximum of 32 characters.

```
OS10(conf-telemetry)# subscription-profile profile-name
```

2 Enter the name of a pre-configured sensor group and sampling interval in SUBSCRIPTION-PROFILE mode. Valid sensor-group names are: `bgp`, `bgp-peer`, `buffer`, `device`, `environment`, `interface`, `lag`, and `system`. To view the data contents of a pre-configured sensor group, use the `show telemetry sensor-group` command. The `interface` sensor group supports only physical and port-channel interfaces.

The sampling interval is in milliseconds, from 0 (whenever an event occurs; near real-time) to 4294967295. The default is 15000. Repeat this step to add sensor groups to the subscription profile.

```
OS10(conf-telemetry-sp-subscription)# sensor-group group-name sampling-interval
```

3 Enter the name of a destination group in SUBSCRIPTION-PROFILE mode. Telemetry data is sent to the IP address and port specified in the destination group. Repeat this step to add destination groups to the subscription profile.

```
OS10(conf-telemetry-sp-subscription)# destination-group name
```

4 Enter the source interface in SUBSCRIPTION-PROFILE mode. The system uses the source interface to derive the VRF instance and IP address used to communicate with destination devices. For gRPC transport, source interface configuration is optional.

```
OS10(conf-telemetry-sp-subscription)# source interface
```

Where *interface* is one of the following values:

- `ethernet node/slot/port[:subport]` — Enter a physical Ethernet interface.
- `loopback number` — Enter a Loopback interface, from 0 to 16383.
- `management 1/1/1` — Enter the management interface.
- `port-channel channel-id` — Enter a port-channel ID, from 1 to 28.
- `vlan vlan-id` — Enter a VLAN ID, from 1 to 4093.

5 Configure the `gpb` encoding format in which data is streamed in SUBSCRIPTION-PROFILE mode.

```
OS10(conf-telemetry-sp-subscription)# encoding format
```

6 Configure the `gRPC` transport protocol used to stream data to a destination in SUBSCRIPTION-PROFILE mode. gRPC with Transport Security Layer (TLS) certificates enabled is the default transport protocol. To disable TLS certificate exchange, use the `transport grpc no-tls` command.

```
OS10(conf-telemetry-sp-subscription)# transport protocol [no-tls]
```

After you configure a subscription profile, the telemetry agent starts collecting data and streaming it to destination devices.

View telemetry configuration

Use the following `show` commands to display telemetry configuration.

```
OS10# show telemetry

Telemetry Status             : enabled

-- Telemetry Destination Groups --
Group : dest1
  Destination : 10.11.56.204  Port : 40001

-- Telemetry Sensor Groups --
Group : bgp
  Sensor Path : bgp/bgp-oper/bgpPrfxCntrsEntry
  Sensor Path : bgp/bgp-oper/bgpPeerCount
Group : bgp-peer
  Sensor Path : infra-bgp/peer-state/peer-status
Group : buffer
  Sensor Path : base-qos/queue-stat
  Sensor Path : base-qos/priority-group-stat
  Sensor Path : base-qos/buffer-pool-stat
  Sensor Path : base-qos/buffer-pool
Group : device
  Sensor Path : base-pas/chassis
  Sensor Path : base-pas/card
  Sensor Path : base-switch/switching-entities/switch-stats
Group : environment
  Sensor Path : base-pas/entity
  Sensor Path : base-pas/psu
  Sensor Path : base-pas/fan-tray
  Sensor Path : base-pas/fan
  Sensor Path : base-pas/led
  Sensor Path : base-pas/temperature
  Sensor Path : base-pas/temp_threshold
  Sensor Path : base-pas/media
  Sensor Path : base-pas/media-channel
Group : interface
  Sensor Path : if/interfaces-state/interface/statistics
  Sensor Path : dell-base-if-cmn/if/interfaces-state/interface
Group : lag
  Sensor Path : dell-base-if-cmn/if/interfaces
Group : system
  Sensor Path : system-status/current-status

-- Telemetry Subscription Profiles --
Name : subscription-1

  Destination Groups(s) : dest1

  Sensor-group        Sample-interval
  ----------------------------------
  bgp                 300000
  bgp-peer            0
  buffer              15000
  device              300000
  environment         300000
  interface           180000
```

```
  lag                 0
  system              300000

  Encoding : gpb
  Transport : grpc    TLS : disabled
  Source Interface : ethernet1/1/1
  Active : true
  Reason : Connection summary: One or more active connections
           The connection 10.11.56.204:40001 is in connected state
```

View destination group

```
OS10# show telemetry destination-group

Telemetry Status             : enabled

-- Telemetry Destination Groups --
Group : dest1
  Destination : 10.11.56.204  Port : 40001
```

View sensor groups

```
OS10# show telemetry sensor-group

Telemetry Status             : enabled

-- Telemetry Sensor Groups --
Group : bgp
  Sensor Path : bgp/bgp-oper/bgpPrfxCntrsEntry
  Sensor Path : bgp/bgp-oper/bgpPeerCount
Group : bgp-peer
  Sensor Path : infra-bgp/peer-state/peer-status
Group : buffer
  Sensor Path : base-qos/queue-stat
  Sensor Path : base-qos/priority-group-stat
  Sensor Path : base-qos/buffer-pool-stat
  Sensor Path : base-qos/buffer-pool
Group : device
  Sensor Path : base-pas/chassis
  Sensor Path : base-pas/card
  Sensor Path : base-switch/switching-entities/switch-stats
Group : environment
  Sensor Path : base-pas/entity
  Sensor Path : base-pas/psu
  Sensor Path : base-pas/fan-tray
  Sensor Path : base-pas/fan
  Sensor Path : base-pas/led
  Sensor Path : base-pas/temperature
  Sensor Path : base-pas/temp_threshold
  Sensor Path : base-pas/media
  Sensor Path : base-pas/media-channel
Group : interface
  Sensor Path : if/interfaces-state/interface/statistics
  Sensor Path : dell-base-if-cmn/if/interfaces-state/interface
Group : lag
  Sensor Path : dell-base-if-cmn/if/interfaces
Group : system
  Sensor Path : system-status/current-status
```

View subscription profiles

```
OS10# show telemetry subscription-profile

Telemetry Status             : enabled

-- Telemetry Subscription Profile --

Name : subscription-1
```

```
  Destination Groups(s) : dest1

  Sensor-group        Sample-interval
  ------------------------------------
  bgp                 300000
  bgp-peer            0
  buffer              15000
  device              300000
  environment         300000
  interface           180000
  lag                 0
  system              300000

  Encoding : gpb
  Transport : grpc    TLS : disabled
  Source Interface : ethernet1/1/1
  Active : true
  Reason : Connection summary: One or more active connections
           The connection 10.11.56.204:40001 is in connected state
```

Verify telemetry in running configuration

```
OS10# show running-configuration telemetry
!
telemetry
 enable
 !
 destination-group dest1
  destination 10.11.56.204 40001
 !
 subscription-profile subscription-1
  destination-group dest1
  sensor-group bgp 300000
  sensor-group bgp-peer 0
  sensor-group buffer 15000
  sensor-group device 300000
  sensor-group environment 300000
  sensor-group interface 180000
  sensor-group lag 0
  sensor-group system 300000
  encoding gpb
  transport grpc no-tls
  source-interface ethernet1/1/1
```

Telemetry commands

debug telemetry

Starts data collection to troubleshoot telemetry operation.

Syntax	`debug telemetry`
Parameters	None
Default	Not configured
Command mode	EXEC
Usage information	When an error condition occurs, use the `debug telemetry` command to store telemetry data in a debug file. The telemetry debug file is stored at `/var/log/grpc_server.log`.
Example	`OS10# debug telemetry`

Supported releases 10.4.3.0 or later

destination

Configures a destination management device that receives streaming telemetry.

Syntax `destination {ip-address | domain-name} port-number`

Parameters

- *ip-address* — Enter the IPv4 or IPv6 address of the destination device. You can enter a fully qualified domain name (FQDN). The destination domain name resolves to an IP address — see System domain name and list.
- *domain-name* — Enter the fully qualified domain name of the destination device. A maximum of 32 characters.
- *port-number* — Enter the transport-service port number to which telemetry data is sent on the destination device.

Default Not configured

Command mode DESTINATION-GROUP

Usage information When you associate a destination group with a subscription, telemetry data is sent to the IP address and port specified by the `destination` command. In the 10.4.3.0 release, only one destination is supported. The `no` version of this command removes the configured destination.

Example

```
OS10(conf-telemetry)# destination-group dest1
OS10(conf-telemetry-dg-dest1)# destination 10.11.56.204 40001
OS10(conf-telemetry-dg-dest1)#
```

Supported releases 10.4.3.0 or later

destination-group (subscription-profile)

Assigns a destination group to a subscription profile for streaming telemetry.

Syntax `destination-group group-name`

Parameters *group-name* — Enter the name of the destination group. A maximum of 32 characters.

Default Not configured

Command mode SUBSCRIPTION-PROFILE

Usage information A subscription profile associates destination groups and sensor groups. A destination group defines the destination servers to which streaming telemetry data is sent. The `no` version of this command removes the configured group from the subscription profile.

Example

```
OS10(conf-telemetry)# subscription-profile subscription-1
OS10(conf-telemetry-sp-subscription-1)# destination-group dest1
```

Supported releases 10.4.3.0 or later

destination-group (telemetry)

Configures a destination group for streaming telemetry.

Syntax	`destination-group` *`group-name`*
Parameters	*`group-name`* — Enter the name of the destination group. A maximum of 32 characters maximum.
Default	Not configured
Command mode	TELEMETRY
Usage information	A destination group defines the destination servers to which streaming telemetry data is sent. The `no` version of this command removes the configured group.
Example	`OS10(conf-telemetry)# destination-group dest1` `OS10(conf-telemetry-dg-dest1)#`
Supported releases	10.4.3.0 or later

enable

Enables telemetry on the switch.

Syntax	`enable`
Parameters	None
Default	Telemetry is disabled.
Command mode	TELEMETRY
Usage information	Enter the `no enable` command to disable telemetry.
Example	`OS10(conf-telemetry)# enable`
Supported releases	10.4.3.0 or later

encoding

Configures the encoding format used to stream telemetry data to a destination device.

Syntax	`encoding` *`format`*
Parameters	*`format`* — Enter the `gpb` (Google protocol buffer) encoding format in which data is streamed.
Default	None
Command mode	SUBSCRIPTION-PROFILE
Usage information	The `no` version of the command removes the configured encoding format from a subscription profile.
Example	`OS10(conf-telemetry)# subscription-profile subscription-1` `OS10(conf-telemetry-sp-subscription-1)# encoding gpb`
Supported releases	10.4.3.0 or later

sensor-group (subscription-profile)

Assigns a sensor group with sampling interval to a subscription profile for streaming telemetry.

Syntax `sensor-group group-name sampling-interval`

Parameters

- *group-name* — Enter the name for the sensor group. In release 10.4.3.0, only pre-configured sensor groups are supported: `bgp`, `bgp-peer`, `buffer`, `device`, `environment`, `interface`, `lag`, and `system`.
- *sampling-interval* — Enter the interval in milliseconds used to collect data samples. The range is 0 to 4294967295. The default is 15000.

Default Not configured

Command mode SUBSCRIPTION-PROFILE

Usage information This command assigns the sensors from which data is collected for streaming telemetry to a subscription profile and specifies the sampling rate. To add sensor groups to the subscription profile, re-enter the command. The `interface` sensor group supports only physical and port-channel interfaces. The `no` version of this command removes the sensor group from the subscription profile.

Example

```
OS10(conf-telemetry)# subscription-profile subscription-1
OS10(conf-telemetry-sp-subscription-1)# sensor-group bgp 30000
OS10(conf-telemetry-sp-subscription-1)# sensor-group environment 415000
```

Supported releases 10.4.3.0 or later

sensor-group (telemetry)

Configures a sensor group for streaming telemetry.

NOTE: This command is not supported in release 10.4.3.0.

Syntax `sensor-group group-name`

Parameters *group-name* — Enter the name of the sensor group. A maximum of 32 characters. You can enter the name of a pre-configured sensor-group profile. Valid values are: `bgp`, `bgp-peer`, `buffer`, `device`, `environment`, `interface`, `lag`, and `system`.

Default Not configured

Command mode TELEMETRY

Usage information A sensor group defines the YANG models from which data is collected and streamed to a destination. You can also use pre-configured sensor group profiles without specifying the YANG container path. In release 10.4.3.0, telemetry supports only pre-configured sensor groups. The `interface` sensor group supports only physical and port-channel interfaces. To view the YANG paths in a pre-configured sensor group, use the `show telemetry sensor-group` command.

Example

```
OS10(conf-telemetry)# sensor-group bgp
OS10(conf-telemetry-sg-bgp)#
```

Supported releases 10.4.3.0 or later

sensor-path

Configures the path to a YANG container from which data is streamed.

NOTE: This command is not supported in release 10.4.3.0.

Syntax `sensor-path yang-path`

Parameters *yang-path* — Enter the YANG path. See YANG-modeled telemetry data for the YANG paths for telemetry data collection.

Default Not configured

Command mode SENSOR-GROUP

Usage information The data in the specified YANG path is streamed to a configured destination at a specified interval or in near real-time when an event occurs. You can enter a sensor path at any level in the container hierarchy. Data from all containers in the hierarchy below this path is streamed to destination devices. To configure multiple sensor paths in a sensor group, re-enter the command.

Example

```
OS10(conf-telemetry)# sensor-group interfaces
OS10(conf-telemetry-sg-interfaces)# sensor-group dell-base-if-cmn/if/
interfaces-state/interface
OS10(conf-telemetry-sg-interfaces)# sensor-group dell-base-cmn/if/interfaces/
interface/member-ports
```

Supported releases 10.4.3.0 or later

show telemetry

Displays the configured destination-group, sensor-group, and subscription profiles for streaming telemetry.

Syntax `show telemetry [destination-group [group-name] | sensor-group [group-name] | subscription-profile [profile-name]]`

Parameters

- `destination-group` — Display only destination groups or a specified group.
- `sensor-group` — Display only sensor groups or a specified group.
- `subscription-profile` — Display only subscription profiles or a specified profile.

Default Display all destination-group, sensor-group, and subscription configurations.

Command mode EXEC

Usage information Use the `show telemetry` command to verify the configured destination devices, sensor data sources, and subscription profiles.

Examples

```
OS10# show telemetry destination-group

Telemetry Status             : enabled

-- Telemetry Destination Groups --
Group : dest1
  Destination : 10.11.56.204  Port : 40001
```

```
OS10# show telemetry sensor-group

Telemetry Status             : enabled
```

```
-- Telemetry Sensor Groups --
Group : bgp
  Sensor Path : bgp/bgp-oper/bgpPrfxCntrsEntry
  Sensor Path : bgp/bgp-oper/bgpPeerCount
Group : bgp-peer
  Sensor Path : infra-bgp/peer-state/peer-status
Group : buffer
  Sensor Path : base-qos/queue-stat
  Sensor Path : base-qos/priority-group-stat
  Sensor Path : base-qos/buffer-pool-stat
  Sensor Path : base-qos/buffer-pool
Group : device
  Sensor Path : base-pas/chassis
  Sensor Path : base-pas/card
  Sensor Path : base-switch/switching-entities/switch-stats
Group : environment
  Sensor Path : base-pas/entity
  Sensor Path : base-pas/psu
  Sensor Path : base-pas/fan-tray
  Sensor Path : base-pas/fan
  Sensor Path : base-pas/led
  Sensor Path : base-pas/temperature
  Sensor Path : base-pas/temp_threshold
  Sensor Path : base-pas/media
  Sensor Path : base-pas/media-channel
Group : interface
  Sensor Path : if/interfaces-state/interface/statistics
  Sensor Path : dell-base-if-cmn/if/interfaces-state/interface
Group : lag
  Sensor Path : dell-base-if-cmn/if/interfaces
Group : system
  Sensor Path : system-status/current-status
```

```
OS10# show telemetry subscription-profile

Telemetry Status             : enabled

-- Telemetry Subscription Profile --

Name : subscription-1

  Destination Groups(s) : dest1

  Sensor-group        Sample-interval
  ------------------------------------
  bgp                 300000
  bgp-peer            0
  buffer              15000
  device              300000
  environment         300000
  interface           180000
  lag                 0
  system              300000

  Encoding : gpb
  Transport : grpc    TLS : disabled
  Source Interface : ethernet1/1/1
  Active : true
  Reason : Connection summary: One or more active connections
           The connection 10.11.56.204:40001 is in connected state
```

Supported releases 10.4.3.0 or later

source interface

Configures the source interface used to stream telemetry data to a destination device.

Syntax `source interface interface`

Parameters *interface* — One of the following values:

- `ethernet node/slot/port[:subport]` — Enter a physical Ethernet interface.
- `loopback number` — Enter a Loopback interface, from 0 to 16383.
- `management 1/1/1` — Enter the management interface.
- `port-channel channel-id` — Enter a port-channel ID, from 1 to 28.
- `vlan vlan-id` — Enter a VLAN ID, from 1 to 4093.

Default None

Command mode SUBSCRIPTION-PROFILE

Usage information The telemetry agent uses the source interface to derive the VRF instance and IP address used to communicate with destination devices. For gRPC transport, source interface configuration is optional. The `no` version of the command removes the configured source interface from a subscription profile.

Example

```
OS10(conf-telemetry)# subscription-profile subscription-1
OS10(conf-telemetry-sp-subscription-1)# source-interface ethernet 1/1/1
```

Supported releases 10.4.3.0 or later

subscription-profile

Configures a subscription profile for streaming telemetry data.

Syntax `subscription-profile profile-name`

Parameters *profile-name* — Enter a profile name. A maximum of 32 characters.

Default Not configured

Command mode TELEMETRY

Usage information A subscription profile associates destination groups with sensor groups, and specifies the data encoding format and transport protocol. Telemetry data is sent to the IP address and port specified in the destination groups.

Example

```
OS10(conf-telemetry)# subscription-profile subscription-1
OS10(conf-telemetry-sp-subscription-1)#
```

Supported releases 10.4.3.0 or later

telemetry

Enters Telemetry configuration mode to configure streaming telemetry.

Syntax `telemetry`

Parameters None

Default	Telemetry is disabled on the switch.
Command mode	CONFIGURATION
Usage information	Enable and disable streaming telemetry in Telemetry mode.
Example	`OS10(config)# telemetry` `OS10(conf-telemetry)#`
Supported releases	10.4.3.0 or later

transport

Configures the transport protocol used to stream telemetry data to a remote management device.

Syntax	`transport` *`protocol`* `[no-tls]`
Parameters	• *`protocol`* — Enter the `gRPC` (Google remote procedure call) transport protocol used for telemetry sessions. • `no-tls` — (Optional) Disable Transport Security Layer (TLS) certificate exchange with gRPC transport.
Default	OS10 telemetry uses the gRPC protocol for transport with TLS certificates enabled.
Command mode	SUBSCRIPTION-PROFILE
Usage information	gRPC with TLS transport is enabled by default. To use gRPC over TLS transport, you must install a X.509v3 certificate on the switch. To disable TLS certificate exchange, use the `transport grpc no-tls` command. **NOTE: gRPC with TLS transport does not support host certificates. To use a CA certificate, see Request and install host certificates.** The `no` version of the command removes the configured transport protocol from a subscription profile.
Example	`OS10(conf-telemetry)# subscription-profile subscription-1` `OS10(conf-telemetry-sp-subscription-1)# transport grpc`
Supported releases	10.4.3.0 or later

Example: Configure streaming telemetry

```
OS10(config)# telemetry
OS10(conf-telemetry)# enable
OS10(conf-telemetry)# destination-group dest1
OS10(conf-telemetry-dg-dest1)# destination 10.11.56.204 40001
OS10(conf-telemetry-dg-dest1)# exit
OS10(conf-telemetry)# subscription-profile subscription-1
OS10(conf-telemetry-sp-subscription-1)# sensor-group bgp 300000
OS10(conf-telemetry-sp-subscription-1)# sensor-group bgp-peer 0
OS10(conf-telemetry-sp-subscription-1)# sensor-group buffer 15000
OS10(conf-telemetry-sp-subscription-1)# sensor-group device 300000
OS10(conf-telemetry-sp-subscription-1)# sensor-group environment 300000
OS10(conf-telemetry-sp-subscription-1)# sensor-group interface 180000
OS10(conf-telemetry-sp-subscription-1)# sensor-group lag 0
OS10(conf-telemetry-sp-subscription-1)# sensor-group system 300000
OS10(conf-telemetry-sp-subscription-1)# destination-group dest1
OS10(conf-telemetry-sp-subscription-1)# encoding gpb
OS10(conf-telemetry-sp-subscription-1)# transport grpc no-tls
```

```
OS10(conf-telemetry-sp-subscription-1)# source-interface ethernet 1/1/1
OS10(conf-telemetry-sp-subscription-1)# end
```

```
OS10# show telemetry

Telemetry Status             : enabled

-- Telemetry Destination Groups --
Group : dest1
  Destination : 10.11.56.204  Port : 40001

-- Telemetry Sensor Groups --
Group : bgp
  Sensor Path : bgp/bgp-oper/bgpPrfxCntrsEntry
  Sensor Path : bgp/bgp-oper/bgpPeerCount
Group : bgp-peer
  Sensor Path : infra-bgp/peer-state/peer-status
Group : buffer
  Sensor Path : base-qos/queue-stat
  Sensor Path : base-qos/priority-group-stat
  Sensor Path : base-qos/buffer-pool-stat
  Sensor Path : base-qos/buffer-pool
Group : device
  Sensor Path : base-pas/chassis
  Sensor Path : base-pas/card
  Sensor Path : base-switch/switching-entities/switch-stats
Group : environment
  Sensor Path : base-pas/entity
  Sensor Path : base-pas/psu
  Sensor Path : base-pas/fan-tray
  Sensor Path : base-pas/fan
  Sensor Path : base-pas/led
  Sensor Path : base-pas/temperature
  Sensor Path : base-pas/temp_threshold
  Sensor Path : base-pas/media
  Sensor Path : base-pas/media-channel
Group : interface
  Sensor Path : if/interfaces-state/interface/statistics
  Sensor Path : dell-base-if-cmn/if/interfaces-state/interface
Group : lag
  Sensor Path : dell-base-if-cmn/if/interfaces
Group : system
  Sensor Path : system-status/current-status

-- Telemetry Subscription Profiles --
Name : subscription-1

  Destination Groups(s) : dest1

  Sensor-group        Sample-interval
  ------------------------------------
  bgp                 300000
  bgp-peer            0
  buffer              15000
  device              300000
  environment         300000
  interface           180000
  lag                 0
  system              300000

  Encoding : gpb
  Transport : grpc    TLS : disabled
  Source Interface : ethernet1/1/1
  Active : true
  Reason : Connection summary: One or more active connections
           The connection 10.11.56.204:40001 is in connected state
```

18

RESTCONF API

RESTCONF is a representational state transfer (REST)-like protocol that uses HTTPS connections. Use the OS10 RESTCONF API to set up the configuration parameters on OS10 switches using JavaScript Object Notation (JSON)-structured messages. Use any programming language to create and send JSON messages. The examples in this chapter use curl.

The OS10 RESTCONF implementation complies with RFC 8040. You can use the RESTCONF API to configure and monitor an OS10 switch.

The OS10 RESTCONF API uses HTTP with the Transport Layer Security (TLS) protocol over port 443. OS10 supports HTTP/1.1 transport as defined in RFC 7230. The RESTCONF API uses pluggable authentication modules (PAM)-based authentication.

On supported platforms, the OS10 RESTCONF API is disabled by default. To configure and enable the RESTCONF API, see the *Configure the RESTCONF API* section.

To configure and monitor an OS10 switch, use REST API client tools, such as Postman or Swagger, to execute web requests. REST API requests, such as GET, PUT, POST, DELETE, and PATCH, operate on OS10 RESTCONF resources, such as:

Table 80. OS10 RESTCONF resources

Resource	Description	URL
Data	Configuration and operational data the RESTCONF API client accesses	/restconf/data
Operations	Container for the protocol-specific data model operations OS10 advertises	/restconf/operations

To browse OS10 RESTCONF API end-points and operations, see the OpenAPI JSON files available on the OS10 Enterprise Edition Software page at the Dell EMC Support site. Download the JSON files and import them to REST API client tools; for example, Swagger or Postman, to generate code, documentation, and test cases. For information about the OpenAPI specification, go to https://swagger.io/docs/specification/about/.

Configure RESTCONF API

To use the RESTCONF API on an OS10 interface, you must enable the RESTCONF API service using the `rest api restconf` command. You can also configure HTTPS access, including:

- Hostname required in a Secure Sockets Layer (SSL) self-signed server certificate
- Timeout for the HTTPS connection
- Cipher suites for encrypting data in an HTTPS connection

After you enable the RESTCONF API, you can send HTTPS requests from a remote device.

1. (Optional) Configure the hostname required in the SSL self-signed server certificate in a RESTCONF HTTPS connection in CONFIGURATION mode, using a maximum of 30 alphanumeric characters. Enter the IP address or domain name of the OS10 switch. By default, the domain name of the OS10 switch is used as the hostname.

   ```
   rest https server-certificate name hostname
   ```

2. (Optional) Configure the timeout that a RESTCONF HTTPS session uses in CONFIGURATION mode, from 30 to 65535 seconds; default 30.

   ```
   rest https session timeout seconds
   ```

3 (Optional) Limit the ciphers that the switch uses in a RESTCONF HTTPS session to encrypt and decrypt data in CONFIGURATION mode. By default, all cipher suites installed on OS10 are supported. Separate multiple entries with a blank space. Valid cipher-suite values are:

- dhe-rsa-with-aes-128-gcm-SHA256
- dhe-rsa-with-aes-256-gcm-SHA384
- ecdhe-rsa-with-aes-128-gcm-SHA256
- ecdhe-rsa-with-aes-256-gcm-SHA384

```
rest https cipher-suite
```

4 Enable RESTCONF API in CONFIGURATION mode.

```
rest api restconf
```

RESTCONF API configuration

```
OS10(config)# rest https server-certificate name OS10.dell.com
OS10(config)# rest https session timeout 60
OS10(config)# rest https cipher-suite dhe-rsa-with-aes-128-gcm-SHA256
dhe-rsa-with-aes-256-gcm-SHA384 ecdhe-rsa-with-aes-256-gcm-SHA384
OS10(config)# rest api restconf
```

CLI commands for RESTCONF API

rest api restconf

Enables the RESTCONF API service on the switch.

Syntax	`rest api restconf`
Parameters	None
Default	RESTCONF API is disabled.
Command Mode	CONFIGURATION
Usage Information	• After you enable the RESTCONF API, you can send curl commands in HTTPS requests from a remote device. • The `no` version of the command disables the RESTCONF API.
Example	`OS10(config)# rest api restconf`
Supported Releases	10.4.1.0 or later

rest https cipher-suite

Limits the ciphers to encrypt and decrypt REST HTTPS data.

Syntax	`rest https cipher-suite` *`cipher-list`*
Parameters	*`cipher-list`* — Enter the ciphers supported in a REST API HTTPS session. Separate multiple entries with a blank space. Valid cipher suites are: • dhe-rsa-with-aes-128-gcm-SHA256 • dhe-rsa-with-aes-256-gcm-SHA384 • ecdhe-rsa-with-aes-128-gcm-SHA256

- ecdhe-rsa-with-aes-256-gcm-SHA384

Default All cipher suites installed with OS10 are supported.

Command Mode CONFIGURATION

Usage Information

- Use the `rest https cipher-suite` command to restrict the ciphers that a RESTCONF HTTPS session uses.
- The `no` version of the command removes the cipher list and restores the default value.

Example

```
OS10(config)# rest https cipher-suite dhe-rsa-with-aes-128-gcm-SHA256
dhe-rsa-with-aes-256-gcm-SHA384 ecdhe-rsa-with-aes-256-gcm-SHA384
```

Supported Releases 10.4.1.0 or later

rest https server-certificate

Creates the SSL self-signed server certificate a RESTCONF HTTPS connection uses.

Syntax `rest https server-certificate name hostname`

Parameters `name hostname` — Enter the IP address or domain name of the OS10 switch.

Default The OS10 switch domain name is used as the *hostname*.

Command Mode CONFIGURATION

Usage Information The `no` version of the command removes the host name from the SSL server certificate.

Example

```
OS10(config)# rest https server-certificate name 10.10.10.10
```

Supported Releases 10.4.1.0 or later

rest https session timeout

Configures the timeout a RESTCONF HTTPS connection uses.

Syntax `rest https session timeout seconds`

Parameters *seconds* — Enter the switch timeout for an HTTPS request from a RESTCONF client, from 30 to 65535 seconds.

Default 30 seconds

Command Mode CONFIGURATION

Usage Information

- If no HTTPS request is received within the configured time, the switch closes the RESTCONF HTTPS session.
- The `no` version of the command removes the configured RESTCONF HTTPS session timeout.

Example

```
OS10# rest https session timeout 60
```

Supported Releases 10.4.1.0 or later

RESTCONF API tasks

Using the RESTCONF API, you can provision OS10 switches using HTTPS requests. The examples in this section show how to access the OS10 RESTCONF API using `curl` commands. curl is a Linux shell command that generates HTTPS requests and is executed on an external server.



curl Commands

curl command options include:

- `-X` specifies the HTTPS request type; for example, `POST` , `PATCH`, or `GET`.
- `-u` specifies the user name and password to use for server authentication.
- `-k` specifies a text file to read curl arguments from. The command line arguments found in the text file will be used as if they were provided on the command line. Use the IP address or URL of the OS10 switch when you access the OS10 RESTCONF API from a remote orchestration system.
- `-H` specifies an extra header to include in the request when sending HTTPS to a server. You can enter multiple extra headers.
- `-d` sends the specified data in an HTTPS request.

In curl commands, use `%2F` to represent a backslash (`/`); for example, enter `ethernet1/2/3` as `ethernet1%2F1%2F3`.

For more information, see the curl Man page.

Usage Information

Consider the following when accessing OS10 RESTSCONF API using `curl` commands:

- The HTTPS request type; `GET EVERYTHING` can cause delay especially in a scaled system. It is recommended to use a specific URI of the target resource to avoid any delay. For example, `curl -X GET -k -u admin:admin https://$TARGET/restconf/data`
- When a RESTCONF query is in progress, you cannot configure any CLI commands until a RESTCONF query is complete.
- It is recommended to use `POST` request instead of `PUT`, to replace the target data resources.

View XML structure of CLI commands

To use the RESTCONF API to configure and monitor an OS10 switch, create an HTTPS request with data parameters in JSON format. The JSON data parameters correspond to the same parameters in the XML structure of an OS10 command.

To display the parameter values in the XML code of an OS10 command as reference, use the `debug cli netconf` command in EXEC mode. In CONFIGURATION mode, use the `do debug cli netconf` command.

This command enables a CLI-to-XML display. At the prompt, enter the OS10 command of the XML request and the reply you need. To exit the CLI-to-XML display, use the `no debug cli netconf` command.

Locate the XML parameters values for the same JSON data arguments. For example, to configure VLAN 20 on an OS10 switch, enter the RESTCONF endpoint and JSON contents in the curl command. Note how the JSON `type` and `name` parameters are displayed in the XML structure of the `interface vlan` command.

- RESTCONF endpoint: `/restconf/data/interfaces`
- JSON data content:

```
{
    "interface": [{
        "type": "iana-if-type:l2vlan",
        "enabled": true,
        "description": "vlan20",
        "name": "vlan20"
    }]
}
```

- curl command:

```
curl -X POST -u admin:admin -k "https://10.11.86.113/restconf/data/interfaces"
-H "accept: application/json" -H "Content-Type: application/json"
-d '{ "interface": [{ "type": "iana-if-type:l2vlan", "enabled": true,
"description":"vlan20",  "name":"vlan20"}]}'
```

To display values for the `type` and `name` parameters in the curl command, display the XML structure of the `interface vlan 20` configuration command:

```
OS10(config)# do debug cli netconf
OS10(config)# interface vlan 10

Request:
<?xml version="1.0" encoding="UTF-8"?>
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <default-operation>merge</default-operation>
    <error-option>stop-on-error</error-option>
    <test-option>set</test-option>
    <config>
      <interfaces xmlns="urn:ietf:params:xml:ns:yang:ietf-interfaces"
xmlns:ianaift="urn:ietf:params:xml:ns:yang:iana-if-type" xmlns:dell-if="http://www.dellemc.com/
networking/os10/dell-interface" xmlns:dell-eth="http://www.dellemc.com/networking/os10/dell-
ethernet" xmlns:dell-lag="http://www.dellemc.com/networking/os10/dell-lag">
        <interface>
          <type>ianaift:l2vlan</type>
          <name>vlan10</name>
        </interface>
      </interfaces>
    </config>
  </edit-config>
</rpc>


Reply:
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="10">
  <ok/>
</rpc-reply>
OS10(config)# do no debug cli netconf
```

RESTCONF API Examples

Some common RESTCONF API operations include configuring system hostname, and interfaces such as loopback interface. The examples in this section use `curl` commands to send the HTTPS request.

System

Configure system hostname

RESTCONF endpoint

```
/restconf/data/dell-system:system/hostname
```

JSON content

```
{
    "hostname":"MyHost"
}
```

Parameters

- `hostname` *`string`* —Enter the hostname of the system. The default is OS10.

Example

```
curl -X PATCH -k -u admin:admin -H "Content-Type: application/json"
https://10.11.86.113/restconf/data/dell-system:system/hostname
-d '{"hostname":"MyHost"}'
```

Interface

Configure a loopback interface

RESTCONF endpoint

```
/restconf/data/interfaces
```

JSON content

```
{
  "interface": [{
    "type": "iana-if-type:softwareLoopback",
    "enabled": true,
    "description":"loopback interface",
    "name":"loopback1"}]
}
```

Parameters

- `type` *`string`* —Enter `iana-if-type:softwareLoopback` for a loopback interface.
- `enabled` *`bool`*— Enter `true` to enable the interface; enter `false` to disable.
- `description` *`string`* — Enter a text string to describe the interface. A maximum of 80 alphanumeric characters.
- `name` *`string`* — Enter `loopback` *`loopback-id`* of the interface, `loopback-id` is from 0 to 16383.

Example

```
curl -X POST  -k -u admin:admin "https://10.11.86.113/restconf/data/
interfaces"
-H "accept: application/json" -H "Content-Type: application/json"
-d '{"interface": [{"type": "iana-if-type:softwareLoopback", "enabled": true,
"description":"loopback interface", "name":"loopback1"}]}'
```

Configure a loopback interface IP address

RESTCONF endpoint

```
/restconf/data/interfaces/interface/loopback1
```

JSON content

```
{
  "dell-ip:ipv4":{
```

```
    "address": {
      "primary-addr":"6.6.6.6/24"
    }
  }
}
```

Parameters

- `primary-addr` *`ip-address/prefix-length`* — Enter the loopback IP address in dotted-decimal A.B.C.D/x format.

Example

```
curl -X POST  -k -u admin:admin "https://10.11.86.113/restconf/data/interfaces/
interface/loopback1"
-H "accept: application/json" -H "Content-Type: application/json"
-d '{"dell-ip:ipv4":{"address": {"primary-addr":"6.6.6.6/24"}}}'
```

19

Troubleshoot OS10

Critical workloads and applications require constant availability. Dell EMC Networking offers tools to help you monitor and troubleshoot problems before they happen.

Packet and flow capture	Manages packet and traffic
Metrics measurement	Pings, round-trip times, jitter, response times, and so on
Analysis and reporting	Metrics and packet capturing
Alerting	Triggers problem reporting
Logging	Captures system history
Performance monitoring	Establishes baselines and defines triggers for detecting performance problems
Mapping and representation	Defines device locations and status

Dell EMC recommends the following best practices:

- View traffic end-to-end from the application's view point.
- Deploy network management infrastructure rapidly, where needed, when needed, and on-demand.
- Extend analysis beyond the network and watch traffic to and from your host.
- Focus on real-time assessment and use trend analysis to backup your conclusions.
- Emphasize *effective* over *absolute* — leverage management solutions that resolve your most common, most expensive problem quickly.
- Address networking performance issues before you focus on the application performance.
- Use methodologies and technologies that fit your network and needs.
- Continuously monitor performance and availability as a baseline for system performance and system up time to quickly separate network issues from application issues.

Diagnostic tools

This section contains information about advanced software and hardware commands to debug, monitor, and troubleshoot network devices.

NOTE: Output examples are for reference purposes only and may not apply to your specific system.

View inventory

Use the `show inventory` command to view the module IDs of the device.

```
OS10# show inventory
Product                : S4048ON
Description            : S4048-ON 48x10GbE, 6x40GbE QSFP+ Interface Module
Software version       : 10.4.3.0
Product Base           :
Product Serial Number  :
Product Part Number    :
```

```
Unit Type                      Part Number  Rev  Piece Part ID             Svc Tag  Exprs Svc
Code
-------------------------------------------------------------------------------------------------
-
* 1  S4048ON                   0J09D3       X01  TW-0J09D3-28298-49Q-0119  FFD7VS1  335 809 304
65
  1  S4048ON-PWR-2-UNKNOWN     0T9FNW       X01  TW-0T9FNW-28298-49Q-0041  AEIOU##  226 457 410
55
  1  S4048ON-FANTRAY-1         0MGDH8       X01  TW-0MGDH8-28298-49Q-0361  AEIOU##  226 457 410
55
  1  S4048ON-FANTRAY-2         0MGDH8       X01  TW-0MGDH8-28298-49Q-0360  AEIOU##  226 457 410
55
  1  S4048ON-FANTRAY-3         0MGDH8       X01  TW-0MGDH8-28298-49Q-0359  AEIOU##  226 457 410
55
```

Boot partition and image

Display system boot partition and image information.

- View all boot information in EXEC mode.
  ```
  show boot
  ```
- View boot details in EXEC mode.
  ```
  show boot detail
  ```

View boot information

```
OS10# show boot
Current system image information:
==================================
Type      Boot Type      Active                  Standby                 Next-Boot
-------------------------------------------------------------------------------------
Node-id 1 Flash Boot     [A] 10.4.3E             [B] 10.4.3E             [A] active
```

View boot detail

```
OS10# show boot detail
Current system image information detail:
=========================================
Type:                     Node-id 1
Boot Type:                Flash Boot
Active Partition:         A
Active SW Version:        10.4.3E
Active SW Build Version:  10.4.3E.85
Active Kernel Version:    Linux 4.9.110
Active Build Date/Time:   2019-02-18T09:06:10Z
Standby Partition:        B
Standby SW Version:       10.4.3E
Standby SW Build Version: 10.4.3E.80
Standby Build Date/Time:  2019-02-17T15:36:08Z
Next-Boot:                active[A]
```

Monitor processes

Display CPU process information.

- View process CPU utilization information in EXEC mode.
- `show processes node-id` *node-id-number* `[pid` *process-id*`]`

View CPU utilization

```
OS10# show processes node-id 1
top - 09:19:32 up 5 days, 6 min,  2 users,  load average: 0.45, 0.39, 0.34
```

```
Tasks: 208 total,   2 running, 204 sleeping,   0 stopped,   2 zombie
%Cpu(s):  9.7 us,  3.9 sy,  0.3 ni, 85.8 id,  0.0 wa,  0.0 hi,  0.3 si,  0.0 st
KiB Mem:   3998588 total,  2089416 used,  1909172 free,   143772 buffers
KiB Swap:   399856 total,        0 used,   399856 free.   483276 cached Mem
  PID USER      PR  NI    VIRT    RES    SHR S  %CPU %MEM     TIME+ COMMAND
    9 root      20   0       0      0      0 S   6.1  0.0   5:22.41 rcuos/1
  819 snmp      20   0   52736   6696   4132 S   6.1  0.2   2:44.18 snmpd
30452 admin     20   0   22076   2524   2100 R   6.1  0.1   0:00.02 top
    1 root      20   0  112100   5840   3032 S   0.0  0.1   0:12.32 systemd
    2 root      20   0       0      0      0 S   0.0  0.0   0:00.00 kthreadd
    3 root      20   0       0      0      0 S   0.0  0.0   0:25.37 ksoftirqd/0
    5 root       0 -20       0      0      0 S   0.0  0.0   0:00.00 kworker/0:+
    7 root      20   0       0      0      0 R   0.0  0.0   5:15.27 rcu_sched
    8 root      20   0       0      0      0 S   0.0  0.0   2:43.64 rcuos/0
   10 root      20   0       0      0      0 S   0.0  0.0   0:00.00 rcu_bh
   11 root      20   0       0      0      0 S   0.0  0.0   0:00.00 rcuob/0
   12 root      20   0       0      0      0 S   0.0  0.0   0:00.00 rcuob/1
   13 root      rt   0       0      0      0 S   0.0  0.0   0:07.30 migration/0
   14 root      rt   0       0      0      0 S   0.0  0.0   0:02.18 watchdog/0
   15 root      rt   0       0      0      0 S   0.0  0.0   0:02.12 watchdog/1
   16 root      rt   0       0      0      0 S   0.0  0.0   0:04.98 migration/1
   17 root      20   0       0      0      0 S   0.0  0.0   0:03.92 ksoftirqd/1
   19 root       0 -20       0      0      0 S   0.0  0.0   0:00.00 kworker/1:+
   20 root       0 -20       0      0      0 S   0.0  0.0   0:00.00 khelper
   21 root      20   0       0      0      0 S   0.0  0.0   0:00.00 kdevtmpfs
   22 root       0 -20       0      0      0 S   0.0  0.0   0:00.00 netns
   23 root      20   0       0      0      0 S   0.0  0.0   0:00.41 khungtaskd
   24 root       0 -20       0      0      0 S   0.0  0.0   0:00.00 writeback
   25 root      25   5       0      0      0 S   0.0  0.0   0:00.00 ksmd
   --more--
```

```
OS10# show processes node-id 1 pid 1019
top - 09:21:58 up 5 days, 8 min,  2 users,  load average: 0.18, 0.30, 0.31
Tasks:   1 total,   0 running,   1 sleeping,   0 stopped,   0 zombie
%Cpu(s):  9.7 us,  3.9 sy,  0.3 ni, 85.8 id,  0.0 wa,  0.0 hi,  0.3 si,  0.0 st
KiB Mem:   3998588 total,  2089040 used,  1909548 free,   143772 buffers
KiB Swap:   399856 total,        0 used,   399856 free.   483276 cached Mem
  PID USER      PR  NI    VIRT    RES    SHR S  %CPU %MEM     TIME+ COMMAND
 1019 root      20   0 1829416 256080  73508 S   6.6  6.4   1212:36 base_nas
OS10#
```

LED settings

Beacon LEDs identify the location of ports and system status with blinking or solid LEDs.

Change current state of the location LED of the system or interface using the following commands:

```
location-led system {node-id | node-id/unit-id} {on | off}
```

```
location-led interface ethernet {chassis/slot/port[:subport]} {on | off}
```

Change the state of system location LED

```
OS10# location-led system 1 on
OS10# location-led system 1 off
```

Change the state of interface location LED

```
OS10# location-led interface ethernet 1/1/1 on
OS10# location-led interface ethernet 1/1/1 off
```

Packet analysis

Use the Linux `tcpdump` command to analyze network packets. Use filters to limit packet collection and output. You must be logged into the Linux shell to use this command. For more information, seeLog into OS10 Device.

Use the Linux `tcpdump` command without parameters to view packets that flow through all interfaces. To write captured packets to a file, use the `-w` parameter. To read the captured file output offline, you can use open source software packages such as wireshark.

Capture packets from Ethernet interface

```
$ tcpdump -i e101-003-0
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on e101-003-0, link-type EN10MB (Ethernet), capture size 262144 bytes
01:39:22.457185 IP 3.3.3.1 > 3.3.3.4: ICMP echo request, id 5320, seq 26, length 64
01:39:22.457281 IP 3.3.3.1 > 3.3.3.4: ICMP echo reply, id 5320, seq 26, length 64
```

Capture two packets from interface

```
$ tcpdump -c 2 -i e101-003-0
listening on e101-003-0, link-type EN10MB (Ethernet), capture size 96 bytes
01:39:22.457185 IP 3.3.3.1 > 3.3.3.4: ICMP echo request, id 5320, seq 26, length 64
01:39:22.457281 IP 3.3.3.1 > 3.3.3.4: ICMP echo reply, id 5320, seq 26, length 64
2 packets captured
13 packets received by filter
0 packets dropped by kernel
```

Capture packets and write to file

```
$ tcpdump -w 06102016.pcap -i e101-003-0
listening on e101-003-0, link-type EN10MB (Ethernet), capture size 96 bytes
32 packets captured
32 packets received by filter
0 packets dropped by kernel
```

Port adapters and modules

Use the `show diag` command to view diagnostics information for OS10 port adapters and hardware modules.

View diagnostic hardware information

```
OS10# show diag
00:00.0 Host bridge: Intel Corporation Atom Processor S1200 Internal (rev 02)
00:01.0 PCI bridge: Intel Corporation Atom Processor S1200 PCI Express Root Port 1 (rev 02)
00:02.0 PCI bridge: Intel Corporation Atom Processor S1200 PCI Express Root Port 2 (rev 02)
00:03.0 PCI bridge: Intel Corporation Atom Processor S1200 PCI Express Root Port 3 (rev 02)
00:04.0 PCI bridge: Intel Corporation Atom Processor S1200 PCI Express Root Port 4 (rev 02)
00:0e.0 IOMMU: Intel Corporation Atom Processor S1200 Internal (rev 02)
00:13.0 System peripheral: Intel Corporation Atom Processor S1200 SMBus 2.0 Controller 0 (rev
02)
00:13.1 System peripheral: Intel Corporation Atom Processor S1200 SMBus 2.0 Controller 1 (rev
02)
00:14.0 Serial controller: Intel Corporation Atom Processor S1200 UART (rev 02)
00:1f.0 ISA bridge: Intel Corporation Atom Processor S1200 Integrated Legacy Bus (rev 02)
01:00.0 Ethernet controller: Broadcom Corporation Device b850 (rev 03)
02:00.0 SATA controller: Marvell Technology Group Ltd. Device 9170 (rev 12)
03:00.0 PCI bridge: Pericom Semiconductor PI7C9X442SL PCI Express Bridge Port (rev 02)
04:01.0 PCI bridge: Pericom Semiconductor PI7C9X442SL PCI Express Bridge Port (rev 02)
04:02.0 PCI bridge: Pericom Semiconductor PI7C9X442SL PCI Express Bridge Port (rev 02)
04:03.0 PCI bridge: Pericom Semiconductor PI7C9X442SL PCI Express Bridge Port (rev 02)
07:00.0 USB controller: Pericom Semiconductor PI7C9X442SL USB OHCI Controller (rev 01)
```

```
07:00.1 USB controller: Pericom Semiconductor PI7C9X442SL USB OHCI Controller (rev 01)
07:00.2 USB controller: Pericom Semiconductor PI7C9X442SL USB EHCI Controller (rev 01)
08:00.0 Ethernet controller: Intel Corporation 82574L Gigabit Network Connection
```

Test network connectivity

Use the `ping` and `traceroute` commands to test network connectivity. When you ping an IP address, you send packets to a destination and wait for a response. If there is no response, the destination is not active. The `ping` command is useful during configuration if you have problems connecting to a hostname or IP address.

When you execute a *traceroute*, the output shows the path a packet takes from your device to the destination IP address. It also lists all intermediate hops (routers) that the packet traverses to reach its destination, including the total number of hops traversed.

Check IPv4 connectivity

```
OS10# ping 172.31.1.255

Type Ctrl-C to abort.

Sending 5, 100-byte ICMP Echos to 172.31.1.255, timeout is 2 seconds:
Reply to request 1 from 172.31.1.208 0 ms
Reply to request 1 from 172.31.1.216 0 ms
Reply to request 1 from 172.31.1.205 16 ms
::
Reply to request 5 from 172.31.1.209 0 ms
Reply to request 5 from 172.31.1.66 0 ms
Reply to request 5 from 172.31.1.87 0 ms
```

Check IPv6 connectivity

```
OS10# ping 100::1

Type Ctrl-C to abort.

Sending 5, 100-byte ICMP Echos to 100::1, timeout is 2 seconds:
!!!!!
Success rate is 100.0 percent (5/5), round-trip min/avg/max = 0/0/0 (ms)
```

Trace IPv4 network route

```
OS10# traceroute www.Dell Networking.com

Translating "www.Dell Networking.com"...domain server (10.11.0.1) [OK]
Type Ctrl-C to abort.

------------------------------------------------
Tracing the route to www.Dell Networking.com (10.11.84.18),
30 hops max, 40 byte packets
------------------------------------------------

TTL Hostname      Probe1     Probe2     Probe3
1   10.11.199.190 001.000 ms 001.000 ms 002.000 ms
2   gwegress-sjc-02.Dell Networking.com (10.11.30.126) 005.000 ms 001.000 ms 001.000 ms
3   fw-sjc-01.Dell Networking.com (10.11.127.254) 000.000 ms 000.000 ms 000.000 ms
4   www.Dell Networking.com (10.11.84.18) 000.000 ms 000.000 ms 000.000 ms
```

Trace IPv6 network route

```
OS10# traceroute 100::1

Type Ctrl-C to abort.

------------------------------------------------
Tracing the route to 100::1, 64 hops max, 60 byte packets
```

```
-----------------------------------------------
Hops Hostname Probe1    Probe2    Probe3
1    100::1 000.000 ms 000.000 ms 000.000 ms

OS10# traceroute 3ffe:501:ffff:100:201:e8ff:fe00:4c8b

Type Ctrl-C to abort.

-----------------------------------------------
Tracing the route to 3ffe:501:ffff:100:201:e8ff:fe00:4c8b,
64 hops max, 60 byte packets
-----------------------------------------------
Hops Hostname  Probe1   Probe2    Probe3
 1  3ffe:501:ffff:100:201:e8ff:fe00:4c8b
             000.000 ms 000.000 ms 000.000 ms
```

View solution ID

Dell EMC networking switches that are part of a larger solution require a solution identifier (ID).

To view the solution ID including the product base, product serial number, and product part number, use the following `show` commands:

View inventory

```
OS10# show inventory
Product          : S6000-ON
Description          : S6000-ON 32x40GbE QSFP+ Interface Module
Software version : 10.4.9999EX
Product Base     :  ECS Gen3
Product Serial Number : APM001123456789
Product Part Number : 900-590-001

Unit Type                 Part Number Rev   Piece Part ID            Svc Tag   Exprs Svc Code
-------------------------------------------------------------------------------------------
 * 1  S4248FB-ON                            CN-0W1K08-77931-647-0017 OS11SIM  539 375 922 22
   1  S4248FB-ON-PWR-2-AC   02RPHX     A00   CN-02RPHX-17972-5BH-00RE
   1  S4248FB-ON-FANTRAY-1  03CH15     A00   CN-03CH15-77931-62T-0039
   1  S4248FB-ON-FANTRAY-2  03CH15     A00   CN-03CH15-77931-62T-0133
   1  S4248FB-ON-FANTRAY-3  03CH15     A00   CN-03CH15-77931-62T-0067
   1  S4248FB-ON-FANTRAY-4  03CH15     A00   CN-03CH15-77931-62T-0034
   1  S4248FB-ON-FANTRAY-5  03CH15     A00   CN-03CH15-77931-62T-0041
```

View license status

```
OS10# show license status
  System Information

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

Vendor Name     :        Dell EMC

Product Name    :        S6000-VM

Hardware Version:

Platform Name   :        x86_64-dell_s6000_vm

PPID            :        VM0S6000006740000ABCD

Service Tag     :        OS11SIM
```

```
Product Base : ECS Gen3

Product Serial Number : APM001123

Product Part Number : 900-590-0
```

View tech-support details

```
OS10# show tech-support

--------------------show inventory------------------------------

Product          : S6000-ON

Description          : S6000-ON 32x40GbE QSFP+ Interface Module

Software version : 10.4.9999EX

Product Base    :  ECS Gen3

Product Serial Number : APM001123456789

Product Part Number : 900-590-001

------------------------------------------------------------------
<<Output Truncated>>
```

View diagnostics

View system diagnostic information using `show` commands. Use the `show hash-algorithm` command to view the current hash algorithms configured for link aggregation group (LAG) and electronic commerce messaging protocol (ECMP).

View environment

```
OS10# show environment

Unit    State             Temperature
-------------------------------------
1       up                43

Thermal sensors
Unit   Sensor-Id        Sensor-name                                  Temperature
-----------------------------------------------------------------------------
1       1           CPU On-Board temp sensor                          32
1       2           Switch board  temp sensor                         28
1       3           System Inlet Ambient-1 temp sensor                27
1       4           System Inlet Ambient-2 temp sensor                25
1       5           System Inlet Ambient-3 temp sensor                26
1       6           Switch board 2 temp sensor                        31
1       7           Switch board 3 temp sensor                        41
1       8           NPU temp sensor                                   43
```

View hash algorithm

```
OS10# show hash-algorithm
   LagAlgo - CRC EcmpAlgo - CRC
```

View inventory

```
OS10# show inventory
Product               : S4048ON
Description           : S4048-ON 48x10GbE, 6x40GbE QSFP+ Interface Module
Software version      : 10.4.3.0
Product Base          :
```

```
Product Serial Number :
Product Part Number   :

Unit Type                    Part Number  Rev  Piece Part ID             Svc Tag  Exprs Svc
Code
--------------------------------------------------------------------------------------------
-
* 1  S4048ON                 0J09D3       X01  TW-0J09D3-28298-49Q-0119  FFD7VS1  335 809 304
65
  1  S4048ON-PWR-2-UNKNOWN    0T9FNW       X01  TW-0T9FNW-28298-49Q-0041  AEIOU##  226 457 410
55
  1  S4048ON-FANTRAY-1        0MGDH8       X01  TW-0MGDH8-28298-49Q-0361  AEIOU##  226 457 410
55
  1  S4048ON-FANTRAY-2        0MGDH8       X01  TW-0MGDH8-28298-49Q-0360  AEIOU##  226 457 410
55
  1  S4048ON-FANTRAY-3        0MGDH8       X01  TW-0MGDH8-28298-49Q-0359  AEIOU##  226 457 410
55
```

View system information

```
OS10# show system

Node Id              : 1
MAC                  : 34:17:eb:f2:9a:c4
Number of MACs       : 256
Up Time              : 2 days 05:57:17

-- Unit 1 --
Status                     : up
System Identifier          : 1
Down Reason                : unknown
Digital Optical Monitoring : disable
System Location LED        : off
Required Type              : S4048
Current Type               : S4048
Hardware Revision          : X01
Software Version           : 10.4.3.0
Physical Ports             : 48x10GbE, 6x40GbE
BIOS                          : 3.21.0.4
System CPLD                   : 15
Master CPLD                   : 12
Slave CPLD                    : 5

-- Power Supplies --
PSU-ID  Status      Type    AirFlow   Fan  Speed(rpm)  Status
-----------------------------------------------------------------
1       not-present
2       up          UNKNOWN NORMAL    1    10704       up


-- Fan Status --
FanTray  Status      AirFlow   Fan  Speed(rpm)  Status
-----------------------------------------------------------------
1        up          NORMAL    1    10108       up
                               2    10069       up

2        up          NORMAL    1    9954        up
                               2    10108       up

3        up          NORMAL    1    9867        up
                               2    9867        up
```

Diagnostic commands

location-led interface

Changes the location LED of the interface.

Syntax

```
location-led interface ethernet {chassis/slot/port[:subport]} {on | off}
```

Parameters

- *chassis/slot/port[:subport]* — Enter the ethernet interface number.
- on | off — Set the interface LED to be on or off.

Default Not configured

Command Mode EXEC

Usage Information Use this command to change the location LED for the specified interface.

Example

```
OS10# location-led interface ethernet 1/1/1 on
OS10# location-led interface ethernet 1/1/1 off
```

Supported Releases 10.3.0E or later

location-led system

Changes the location LED of the system.

Syntax

```
location-led system {node-id | node-id/unit-id} {on | off}
```

Parameters

- *node-id | node-id/unit-id* — Enter the system ID.
- on | off — Set the system LED to be on or off.

Default Not configured

Command Mode EXEC

Usage Information Use this command to change the location LED for the specified system ID.

Example

```
OS10# location-led system 1 on
OS10# location-led system 1 off
```

Supported Releases 10.3.0E or later

ping

Tests network connectivity to an IPv4 device.

Syntax

```
ping [vrf {management | vrf-name}] [-aAbBdDfhLnOqrRUvV] [-c count] [-i
interval] [-I interface] [-m mark] [-M pmtudisc_option] [-l preload] [-p
pattern] [-Q tos] [-s packetsize] [-S sndbuf] [-t ttl] [-T timestamp_option] [-
w deadline] [-W timeout] [hop1 ...] destination
```

Parameters

- vrf management — (Optional) Pings an IPv4 address in the management virtual routing and forwarding (VRF) instance.
- vrf *vrf-name* — (Optional) Ping an IP address in a specified VRF instance.
- -a — (Optional) Audible ping.

- `-A` — (Optional) Adaptive ping. An inter-packet interval adapts to the round-trip time so that one (or more, if you set the preload option) unanswered probe is present in the network. The minimum interval is 200 msec for a non-super user, which corresponds to Flood mode on a network with a low round-trip time.
- `-b` — (Optional) Pings a broadcast address.
- `-B` — (Optional) Does not allow ping to change the source address of probes. The source address is bound to the address used when the ping starts.
- `-c count` — (Optional) Stops the ping after sending the specified number of ECHO_REQUEST packets until the timeout expires.
- `-d` — (Optional) Sets the SO_DEBUG option on the socket being used.
- `-D` — (Optional) Prints the timestamp before each line.
- `-h` — (Optional) Displays help for this command.
- `-i interval` — (Optional) Enter the interval in seconds to wait between sending each packet, the default is 1 second.
- `-i` *interval*— (Optional) Enter the number of seconds to wait before sending the next packet, from 0 to 60, default 1.
- `-I interface-address` — (Optional) Enter the source interface address with no spaces:
  - For a physical Ethernet interface, enter `ethernet`*`node/slot/port`*; for example, `ethernet1/1/1`.
  - For a VLAN interface, enter `vlan`*`vlan-id`*; for example, `vlan10`.
  - For a Loopback interface, enter `loopback`*`id`*; for example, `loopback1`.
  - For a port-channel interface, enter `port-channel`*`channel-id`*; for example, `port-channel`.
- `-l preload` — (Optional) Enter the number of packets that ping sends before waiting for a reply. Only a super user may preload more than three.
- `-L` — (Optional) Suppress the Loopback of multicast packets for a multicast target address.
- `-m mark` — (Optional) Tags the packets sent to ping a remote device. Use this option with policy routing.
- `-M pmtudisc_option` — (Optional) Enter the path MTU (PMTU) discovery strategy:
  - `do` prevents fragmentation, including local.
  - `want` performs PMTU discovery and fragments large packets locally.
  - `dont` does not set the Don't Fragment (DF) flag.
- `-p pattern` — (Optional) Enter a maximum of 16 pad bytes to fill out the packet you send to diagnose data-related problems in the network; for example, `-p ff` fills the sent packet with all 1's.
- `-Q tos` — (Optional) Enter a maximum of 1500 bytes in decimal or hex datagrams to set quality of service (QoS)-related bits.
- `-s packetsize` — (Optional) Enter the number of data bytes to send, from 1 to 65468, default 56.
- `-S sndbuf` — (Optional) Set the sndbuf socket. By default, the sndbuf socket buffers one packet maximum.
- `-t ttl` — (Optional) Enter the IPv4 time-to-live (TTL) value in seconds.
- `-T timestamp option` — (Optional) Set special IP timestamp options. Valid values for *timestamp option* — `tsonly` (only timestamps), `tsandaddr` (timestamps and addresses), or `tsprespec host1 [host2 [host3 [host4]]]` (timestamp pre-specified hops).
- `-v` — (Optional) Verbose output.
- `-V` — (Optional) Display the version and exit.
- `-w deadline` — (Optional) Enter the time-out value in seconds before the ping exits regardless of how many packets send or receive.
- `-W timeout` — (Optional) Enter the time to wait for a response in seconds. This setting affects the time-out only if there is no response, otherwise ping waits for two round-trip times (RTTs).
- `hop1 ...` (Optional) Enter the IPv4 addresses of the pre-specified hops for the ping packet to take.
- `target` — Enter the IP address you are testing connectivity on.

Default Not configured

Command Mode EXEC

Usage Information

This command uses an ICMP ECHO_REQUEST datagram to receive an ICMP ECHO_RESPONSE from a network host or gateway. Each ping packet has an IPv4 and ICMP header, then a time value and a number of "pad" bytes used to fill out the packet. A ping operation sends a packet to a specified IP address and then measures the time it takes to get a response from the address or device.

If the destination IP address is active, replies are sent back from the server including the IP address, number of bytes sent, lapse time in milliseconds, and TTL, which is the number of hops back from the source to the destination.

Example

```
OS10# ping 20.1.1.1
PING 20.1.1.1 (20.1.1.1) 56(84) bytes of data.
64 bytes from 20.1.1.1: icmp_seq=1 ttl=64 time=0.079 ms
64 bytes from 20.1.1.1: icmp_seq=2 ttl=64 time=0.081 ms
64 bytes from 20.1.1.1: icmp_seq=3 ttl=64 time=0.133 ms
64 bytes from 20.1.1.1: icmp_seq=4 ttl=64 time=0.124 ms
^C
--- 20.1.1.1 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 2997ms
rtt min/avg/max/mdev = 0.079/0.104/0.133/0.025 ms
```

Supported Releases 10.2.0E or later

ping6

Tests network connectivity to an IPv6 device.

Syntax

```
ping6 [vrf {management | vrf-name}] [-aAbBdDfhLnOqrRUvV] [-c count] [-i
interval] [-I interface] [-l preload] [-m mark] [-M pmtudisc_option] [-N
nodeinfo_option] [-p pattern] [-Q tclass] [-s packetsize] [-S sndbuf] [-t ttl]
[-T timestamp_option] [-w deadline] [-W timeout] destination
```

Parameters

- `vrf management` — (Optional) Pings an IPv6 address in the management VRF instance.
- `vrf vrf-name` — (Optional) Pings an IPv6 address in a specified VRF instance.
- `-a` — (Optional) Audible ping.
- `-A` — (Optional) Adaptive ping. An inter-packet interval adapts to the round-trip time so that one (or more, if you set the preload option) unanswered probe is present in the network. The minimum interval is 200 msec for a non-super user, which corresponds to Flood mode on a network with a low round-trip time.
- `-b` — (Optional) Pings a broadcast address.
- `-B` — (Optional) Does not allow ping to change the source address of probes. The source address is bound to the address used when the ping starts.
- `-c count` — (Optional) Stops the ping after sending the specified number of ECHO_REQUEST packets until the timeout expires.
- `-d` — (Optional) Sets the SO_DEBUG option on the socket being used.
- `-D` — (Optional) Prints the timestamp before each line.
- `-F flowlabel` — (Optional) Sets a 20-bit flow label on echo request packets. If value is zero, the kernel allocates a random flow label.
- `-h` — (Optional) Displays help for this command.
- `-i` *interval*— (Optional) Enter the number of seconds to wait before sending the next packet, from 0 to 60, default 1.
- `-i interval` — (Optional) Enter the interval in seconds to wait between sending each packet, the default is 1 second.
- `-I interface-address` — (Optional) Enter the source interface address with no spaces:
  - For a physical Ethernet interface, enter `ethernet`*node/slot/port*; for example, `ethernet1/1/1`.
  - For a VLAN interface, enter `vlan`*vlan-id*; for example, `vlan10`.

  - For a Loopback interface, enter `loopback`*`id`*; for example, `loopback1`.
  - For a port-channel interface, enter `port-channel`*`channel-id`*; for example, `port-channel`.
- `-l` *`preload`* — (Optional) Enter the number of packets that ping sends before waiting for a reply. Only a super-user may preload more than three.
- `-L` — (Optional) Suppress the Loopback of multicast packets for a multicast target address.
- `-m mark` — (Optional) Tags the packets sent to ping a remote device. Use this option with policy routing.
- `-M` *`pmtudisc_option`* — (Optional) Enter the path MTU (PMTU) discovery strategy:
  - `do` prevents fragmentation, including local.
  - `want` performs PMTU discovery and fragments large packets locally.
  - `dont` does not set the Don't Fragment (DF) flag.
- `-p` *`pattern`* — (Optional) Enter a maximum of 16 pad bytes to fill out the packet you send to diagnose data-related problems in the network; for example, `-p ff` fills the sent packet with all 1's.
- `-Q` *`tos`* — (Optional) Enter a maximum of 1500 bytes in decimal or hex datagrams to set the quality of service (QoS)-related bits.
- `-s` *`packetsize`* — (Optional) Enter the number of data bytes to send, from 1 to 65468, default 56.
- `-S` *`sndbuf`* — (Optional) Set the sndbuf socket. By default, the sndbuf socket buffers one packet maximum.
- `-t` *`ttl`* — (Optional) Enter the IPv6 time-to-live (TTL) value in seconds.
- `-T` *`timestamp option`* — (Optional) Set special IP timestamp options. Valid values for *timestamp option* — `tsonly` (only timestamps), `tsandaddr` (timestamps and addresses), or `tsprespec` *`host1 [host2 [host3 [host4]]]`* (timestamp pre-specified hops).
- `-v` — (Optional) Verbose output.
- `-V` — (Optional) Display the version and exit.
- `-w` *`deadline`* — (Optional) Enter the time-out value in seconds before the ping exits regardless of how many packets are sent or received.
- `-W` *`timeout`* — (Optional) Enter the time to wait for a response in seconds. This setting affects the time-out only if there is no response, otherwise ping waits for two round-trip times (RTTs).
- *`hop1 ...`* (Optional) Enter the IPv6 addresses of the pre-specified hops for the ping packet to take.
- *`target`* — Enter the IPv6 destination address in A:B::C:D format, where you are testing connectivity.

Default Not configured

Command Mode EXEC

Usage Information This command uses an ICMP ECHO_REQUEST datagram to receive an ICMP ECHO_RESPONSE from a network host or gateway. Each ping packet has an IPv6 and ICMP header, then a time value and a number of "pad" bytes used to fill out the packet. A pingv6 operation sends a packet to a specified IPv6 address and then measures the time it takes to get a response from the address or device.

Example

```
OS10# ping6 20::1
PING 20::1(20::1) 56 data bytes
64 bytes from 20::1: icmp_seq=1 ttl=64 time=2.07 ms
64 bytes from 20::1: icmp_seq=2 ttl=64 time=2.21 ms
64 bytes from 20::1: icmp_seq=3 ttl=64 time=2.37 ms
64 bytes from 20::1: icmp_seq=4 ttl=64 time=2.10 ms
^C
--- 20::1 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3005ms
rtt min/avg/max/mdev = 2.078/2.194/2.379/0.127 ms
```

Supported Releases 10.2.0E or later

show boot

Displays boot partition-related information.

Syntax `show boot [detail]`

Parameters `detail` — (Optional) Enter to display detailed information.

Default Not configured

Command Mode EXEC

Usage Information Use the `boot system` command to set the boot partition for the next reboot.

Example

```
OS10# show boot
Current system image information:
==================================
Type      Boot Type      Active                   Standby                  Next-Boot
-------------------------------------------------------------------------------------
Node-id 1 Flash Boot     [A] 10.4.3E             [B] 10.4.3E              [A] active
```

Example (Detail)

```
OS10# show boot detail
Current system image information detail:
=========================================
Type:                      Node-id 1
Boot Type:                 Flash Boot
Active Partition:          A
Active SW Version:         10.4.3E
Active SW Build Version:   10.4.3E.85
Active Kernel Version:     Linux 4.9.110
Active Build Date/Time:    2019-02-18T09:06:10Z
Standby Partition:         B
Standby SW Version:        10.4.3E
Standby SW Build Version:  10.4.3E.80
Standby Build Date/Time:   2019-02-17T15:36:08Z
Next-Boot:                 active[A]
```

Supported Releases 10.2.0E or later

show diag

Displays diagnostic information for port adapters and modules.

Syntax `show diag`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show diag
00:00.0 Host bridge: Intel Corporation Atom processor C2000 SoC Transaction
Router (rev 02)
00:01.0 PCI bridge: Intel Corporation Atom processor C2000 PCIe Root Port 1
(rev 02)
00:02.0 PCI bridge: Intel Corporation Atom processor C2000 PCIe Root Port 2
(rev 02)
00:03.0 PCI bridge: Intel Corporation Atom processor C2000 PCIe Root Port 3
(rev 02)
00:04.0 PCI bridge: Intel Corporation Atom processor C2000 PCIe Root Port 4
```

```
(rev 02)
00:0e.0 Host bridge: Intel Corporation Atom processor C2000 RAS (rev 02)
00:0f.0 IOMMU: Intel Corporation Atom processor C2000 RCEC (rev 02)
00:13.0 System peripheral: Intel Corporation Atom processor C2000 SMBus 2.0
(rev 02)
00:14.0 Ethernet controller: Intel Corporation Ethernet Connection I354 (rev
03)
00:14.1 Ethernet controller: Intel Corporation Ethernet Connection I354 (rev
03)
00:16.0 USB controller: Intel Corporation Atom processor C2000 USB Enhanced
Host Controller (rev 02)
00:17.0 SATA controller: Intel Corporation Atom processor C2000 AHCI SATA2
Controller (rev 02)
00:18.0 SATA controller: Intel Corporation Atom processor C2000 AHCI SATA3
Controller (rev 02)
00:1f.0 ISA bridge: Intel Corporation Atom processor C2000 PCU (rev 02)
00:1f.3 SMBus: Intel Corporation Atom processor C2000 PCU SMBus (rev 02)
01:00.0 Ethernet controller: Broadcom Corporation Device b340 (rev 01)
01:00.1 Ethernet controller: Broadcom Corporation Device b340 (rev 01)
```

Supported Releases 10.2.0E or later

show environment

Displays information about environmental system components, such as temperature, fan, and voltage.

Syntax `show environment`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show environment

Unit    State             Temperature
-------------------------------------
1       up                43

Thermal sensors
Unit   Sensor-Id        Sensor-name                                  Temperature
-------------------------------------------------------------------------------
1       1           CPU On-Board temp sensor                          32
1       2           Switch board  temp sensor                         28
1       3           System Inlet Ambient-1 temp sensor                27
1       4           System Inlet Ambient-2 temp sensor                25
1       5           System Inlet Ambient-3 temp sensor                26
1       6           Switch board 2 temp sensor                        31
1       7           Switch board 3 temp sensor                        41
1       8           NPU temp sensor                                   43
```

Supported Releases 10.2.0E or later

show hash-algorithm

Displays hash algorithm information.

Syntax `show hash-algorithm`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show hash-algorithm
   LagAlgo - CRC EcmpAlgo - CRC
```

Supported Releases 10.2.0E or later

show inventory

Displays system inventory information.

Syntax `show inventory`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show inventory
Product                : S4048ON
Description            : S4048-ON 48x10GbE, 6x40GbE QSFP+ Interface Module
Software version       : 10.4.3.0
Product Base           :
Product Serial Number  :
Product Part Number    :

Unit Type                      Part Number  Rev  Piece Part ID              Svc Tag  Expr
-------------------------------------------------------------------------------------------
* 1  S4048ON                   0J09D3       X01  TW-0J09D3-28298-49Q-0119  FFD7VS1  335
  1  S4048ON-PWR-2-UNKNOWN     0T9FNW       X01  TW-0T9FNW-28298-49Q-0041  AEIOU##  226
  1  S4048ON-FANTRAY-1         0MGDH8       X01  TW-0MGDH8-28298-49Q-0361  AEIOU##  226
  1  S4048ON-FANTRAY-2         0MGDH8       X01  TW-0MGDH8-28298-49Q-0360  AEIOU##  226
  1  S4048ON-FANTRAY-3         0MGDH8       X01  TW-0MGDH8-28298-49Q-0359  AEIOU##  226
```

Supported Releases 10.2.0E or later

show processes

View process CPU utilization information.

Syntax `show processes node-id` *`node-id-number`* `[pid` *`process-id`*`]`

Parameters

- *`node-id-number`* — Enter the Node ID number as 1.
- *`process-id`* — (Optional) Enter the process ID number, from 1 to 2147483647.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show processes node-id 1
top - 09:19:32 up 5 days, 6 min,  2 users,  load average: 0.45, 0.39, 0.34
Tasks: 208 total,   2 running, 204 sleeping,   0 stopped,   2 zombie
```

```
%Cpu(s):  9.7 us,  3.9 sy,  0.3 ni, 85.8 id,  0.0 wa,  0.0 hi,  0.3 si,  0.0 st
KiB Mem:   3998588 total,  2089416 used,  1909172 free,   143772 buffers
KiB Swap:   399856 total,        0 used,   399856 free.   483276 cached Mem
  PID USER      PR  NI    VIRT    RES    SHR S  %CPU %MEM     TIME+ COMMAND
    9 root      20   0       0      0      0 S   6.1  0.0   5:22.41 rcuos/1
  819 snmp      20   0   52736   6696   4132 S   6.1  0.2   2:44.18 snmpd
30452 admin     20   0   22076   2524   2100 R   6.1  0.1   0:00.02 top
    1 root      20   0  112100   5840   3032 S   0.0  0.1   0:12.32 systemd
    2 root      20   0       0      0      0 S   0.0  0.0   0:00.00 kthreadd
    3 root      20   0       0      0      0 S   0.0  0.0   0:25.37 ksoftirqd/0
    5 root       0 -20       0      0      0 S   0.0  0.0   0:00.00 kworker/0:+
    7 root      20   0       0      0      0 R   0.0  0.0   5:15.27 rcu_sched
    8 root      20   0       0      0      0 S   0.0  0.0   2:43.64 rcuos/0
   10 root      20   0       0      0      0 S   0.0  0.0   0:00.00 rcu_bh
   11 root      20   0       0      0      0 S   0.0  0.0   0:00.00 rcuob/0
   12 root      20   0       0      0      0 S   0.0  0.0   0:00.00 rcuob/1
   13 root      rt   0       0      0      0 S   0.0  0.0   0:07.30 migration/0
   14 root      rt   0       0      0      0 S   0.0  0.0   0:02.18 watchdog/0
   15 root      rt   0       0      0      0 S   0.0  0.0   0:02.12 watchdog/1
   16 root      rt   0       0      0      0 S   0.0  0.0   0:04.98 migration/1
   17 root      20   0       0      0      0 S   0.0  0.0   0:03.92 ksoftirqd/1
   19 root       0 -20       0      0      0 S   0.0  0.0   0:00.00 kworker/1:+
   20 root       0 -20       0      0      0 S   0.0  0.0   0:00.00 khelper
   21 root      20   0       0      0      0 S   0.0  0.0   0:00.00 kdevtmpfs
   22 root       0 -20       0      0      0 S   0.0  0.0   0:00.00 netns
   23 root      20   0       0      0      0 S   0.0  0.0   0:00.41 khungtaskd
   24 root       0 -20       0      0      0 S   0.0  0.0   0:00.00 writeback
   25 root      25   5       0      0      0 S   0.0  0.0   0:00.00 ksmd
   --more--
```

```
OS10# show processes node-id 1 pid 1019
top - 09:21:58 up 5 days, 8 min,  2 users,  load average: 0.18, 0.30, 0.31
Tasks:   1 total,   0 running,   1 sleeping,   0 stopped,   0 zombie
%Cpu(s):  9.7 us,  3.9 sy,  0.3 ni, 85.8 id,  0.0 wa,  0.0 hi,  0.3 si,  0.0 st
KiB Mem:   3998588 total,  2089040 used,  1909548 free,   143772 buffers
KiB Swap:   399856 total,        0 used,   399856 free.   483276 cached Mem
  PID USER      PR  NI    VIRT    RES    SHR S  %CPU %MEM     TIME+ COMMAND
 1019 root      20   0 1829416 256080  73508 S   6.6  6.4   1212:36 base_nas
OS10#
```

Supported Releases 10.3.0E or later

show system

Displays system information.

Syntax `show system [brief | node-id]`

Parameters

- `brief` — View an abbreviated list of the system information.
- `node-id` — View the node ID number.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show system

Node Id              : 1
MAC                  : 34:17:eb:f2:9a:c4
Number of MACs       : 256
Up Time              : 2 days 05:57:17
```

```
-- Unit 1 --
Status                     : up
System Identifier          : 1
Down Reason                : unknown
Digital Optical Monitoring : disable
System Location LED        : off
Required Type              : S4048
Current Type               : S4048
Hardware Revision          : X01
Software Version           : 10.4.3.0
Physical Ports             : 48x10GbE, 6x40GbE
BIOS                          : 3.21.0.4
System CPLD                   : 15
Master CPLD                   : 12
Slave CPLD                    : 5

-- Power Supplies --
PSU-ID  Status      Type    AirFlow   Fan  Speed(rpm)  Status
-------------------------------------------------------------
1       not-present
2       up          UNKNOWN NORMAL    1    10704       up


-- Fan Status --
FanTray  Status      AirFlow   Fan  Speed(rpm)  Status
-------------------------------------------------------------
1        up          NORMAL    1    10108       up
                               2    10069       up

2        up          NORMAL    1    9954        up
                               2    10108       up

3        up          NORMAL    1    9867        up
                               2    9867        up
```

Example (node-id)

```
OS10# show system node-id 1 fanout-configured

Interface      Breakout capable     Breakout state
-----------------------------------------------------
Eth 1/1/1         No                BREAKOUT_1x1
Eth 1/1/2         No                BREAKOUT_1x1
Eth 1/1/3         No                BREAKOUT_1x1
Eth 1/1/4         No                BREAKOUT_1x1
Eth 1/1/5         No                BREAKOUT_1x1
Eth 1/1/6         No                BREAKOUT_1x1
Eth 1/1/7         No                BREAKOUT_1x1
Eth 1/1/8         No                BREAKOUT_1x1
Eth 1/1/9         No                BREAKOUT_1x1
Eth 1/1/10        No                BREAKOUT_1x1
Eth 1/1/11        No                BREAKOUT_1x1
Eth 1/1/12        No                BREAKOUT_1x1
Eth 1/1/13        No                BREAKOUT_1x1
Eth 1/1/14        No                BREAKOUT_1x1
Eth 1/1/15        No                BREAKOUT_1x1
Eth 1/1/16        No                BREAKOUT_1x1
Eth 1/1/17        No                BREAKOUT_1x1
Eth 1/1/18        No                BREAKOUT_1x1
Eth 1/1/19        No                BREAKOUT_1x1
Eth 1/1/20        No                BREAKOUT_1x1
Eth 1/1/21        No                BREAKOUT_1x1
Eth 1/1/22        No                BREAKOUT_1x1
Eth 1/1/23        No                BREAKOUT_1x1
Eth 1/1/24        No                BREAKOUT_1x1
Eth 1/1/25        No                BREAKOUT_1x1
Eth 1/1/26        No                BREAKOUT_1x1
Eth 1/1/27        No                BREAKOUT_1x1
Eth 1/1/28        No                BREAKOUT_1x1
Eth 1/1/29        No                BREAKOUT_1x1
```

```
Eth 1/1/30          No                BREAKOUT_1x1
Eth 1/1/31          No                BREAKOUT_1x1
Eth 1/1/32          No                BREAKOUT_1x1
Eth 1/1/33          No                BREAKOUT_1x1
Eth 1/1/34          No                BREAKOUT_1x1
Eth 1/1/35          No                BREAKOUT_1x1
Eth 1/1/36          No                BREAKOUT_1x1
Eth 1/1/37          No                BREAKOUT_1x1
Eth 1/1/38          No                BREAKOUT_1x1
Eth 1/1/39          No                BREAKOUT_1x1
Eth 1/1/40          No                BREAKOUT_1x1
Eth 1/1/41          No                BREAKOUT_1x1
Eth 1/1/42          No                BREAKOUT_1x1
Eth 1/1/43          No                BREAKOUT_1x1
Eth 1/1/44          No                BREAKOUT_1x1
Eth 1/1/45          No                BREAKOUT_1x1
Eth 1/1/46          No                BREAKOUT_1x1
Eth 1/1/47          No                BREAKOUT_1x1
Eth 1/1/48          No                BREAKOUT_1x1
Eth 1/1/49          Yes               BREAKOUT_1x1
Eth 1/1/50          Yes               BREAKOUT_1x1
Eth 1/1/51          Yes               BREAKOUT_1x1
Eth 1/1/52          Yes               BREAKOUT_1x1
Eth 1/1/53          Yes               BREAKOUT_1x1
Eth 1/1/54          Yes               BREAKOUT_1x1
```

Example (brief)

```
OS10# show system brief

Node Id           : 1
MAC               : 34:17:eb:f2:9a:c4

-- Unit --
Unit  Status      ReqType     CurType     Version
-----------------------------------------------------------------
1     up          S4048       S4048       10.4.3E

-- Power Supplies --
PSU-ID  Status      Type    AirFlow   Fan  Speed(rpm)  Status
-----------------------------------------------------------------
1       not-present
2       up          UNKNOWN NORMAL    1    10704       up


-- Fan Status --
FanTray  Status      AirFlow   Fan  Speed(rpm)  Status
-----------------------------------------------------------------
1        up          NORMAL    1    9929        up
                               2    9980        up

2        up          NORMAL    1    10095       up
                               2    10082       up

3        up          NORMAL    1    9867        up
                               2    10173       up
```

Supported Releases 10.2.0E or later

traceroute

Displays the routes that packets take to travel to an IP address.

Syntax

```
traceroute [vrf {management | vrf-name}] host [-46dFITnreAUDV] [-f first_ttl]
[-g gate,...] [-i device] [-m max_ttl] [-N squeries] [-p port] [-t tos] [-l
```

```
flow_label] [-w waittime] [-q nqueries] [-s src_addr] [-z sendwait] [--
fwmark=num] host [packetlen]
```

Parameters

- `vrf management` — (Optional) Traces the route to an IP address in the management VRF instance.
- `vrf vrf-name` — (Optional) Traces the route to an IP address in the specified VRF instance.
- *host* — Enter the host to trace packets from.
- `-i interface` — (Optional) Enter the IP address of the interface through which traceroute sends packets. By default, the interface is selected according to the routing table.
- `-m max_ttl` — (Optional) Enter the maximum number of hops, the maximum time-to-live value, that traceroute probes. The default is 30.
- `-p port` — (Optional) Enter a destination port:
  - For UDP tracing, enter the destination port base that traceroute uses. The destination port number is incremented by each probe.
  - For Internet Control Message Protocol (ICMP) tracing, enter the initial ICMP sequence value, incremented by each probe.
  - For TCP tracing, enter the constant destination port to connect.
  - `-P protocol` — (Optional) Use a raw packet of the specified protocol for traceroute. The default protocol is 253 (RFC 3692).
  - `-s source_address` — (Optional) Enter an alternative source address of one of the interfaces. By default, the address of the outgoing interface is used.
  - `-q nqueries` — (Optional) Enter the number of probe packets per hop. The default is 3.
  - `-N squeries` — (Optional) Enter the number of probe packets sent out simultaneously to accelerate traceroute. The default is 16.
  - `-t tos` — (Optional) For IPv4, enter the type of service (ToS) and precedence values to use. `16` sets a low delay; `8` sets a high throughput.
  - `-UL` — (Optional) Use UDPLITE for tracerouting. The default port is 53.
  - `-w waittime` — (Optional) Enter the time in seconds to wait for a response to a probe. The default is 5 seconds.
  - `-z sendwait` — (Optional) Enter the minimal time interval to wait between probes. The default is 0. A value greater than 10 specifies a number in milliseconds, otherwise it specifies a number of seconds. This option is useful when routers rate-limit ICMP messages.
  - `--mtu` — (Optional) Discovers the maximum transmission unit (MTU) from the path being traced.
  - `--back` — (Optional) Prints the number of backward hops when different from the forward direction.
  - *host* — (Required) Enter the name or IP address of the destination device.
  - *packet_len* — (Optional) Enter the total size of the probing packet. The default is 60 bytes for IPv4 and 80 for IPv6.

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# traceroute www.dell.com
traceroute to www.dell.com (23.73.112.54), 30 hops max, 60 byte packets
 1  10.11.97.254 (10.11.97.254)  4.298 ms  4.417 ms  4.398 ms
 2  10.11.3.254 (10.11.3.254)  2.121 ms  2.326 ms  2.550 ms
 3  10.11.27.254 (10.11.27.254)  2.233 ms  2.207 ms  2.391 ms
 4  Host65.hbms.com (63.80.56.65)  3.583 ms  3.776 ms  3.757 ms
 5  host33.30.198.65 (65.198.30.33)  3.758 ms  4.286 ms  4.221 ms
 6  3.GigabitEthernet3-3.GW3.SCL2.ALTER.NET (152.179.99.173)  4.428 ms  2.593
ms  3.243 ms
 7  0.xe-7-0-1.XL3.SJC7.ALTER.NET (152.63.48.254)  3.915 ms  3.603 ms  3.790 ms
 8  TenGigE0-4-0-5.GW6.SJC7.ALTER.NET (152.63.49.254)  11.781 ms  10.600 ms
```

```
9.402 ms
 9  23.73.112.54 (23.73.112.54)  3.606 ms  3.542 ms  3.773 ms
```

Example (IPv6)

```
OS10# traceroute 20::1
traceroute to 20::1 (20::1), 30 hops max, 80 byte packets
 1  20::1 (20::1)  2.622 ms  2.649 ms  2.964 ms
```

Supported Releases 10.2.0E or later

Password recovery

You may need to recover a lost password.

1. Connect to the serial console port. The serial settings are 115200 baud, 8 data bits, and no parity.
2. Reboot or power up the system.
3. Press **ESC** at the Grub prompt to view the boot menu. The OS10-A partition is selected by default.

```
+-----------------------------------------+
|*OS10-A                                  |
| OS10-B                                  |
| ONIE                                    |
+-----------------------------------------+
```

4. Press **e** to open the OS10 GRUB editor.
5. Use the arrow keys to highlight the line that starts with `linux`. Add `init=bin/bash` at the end of the line.

```
+-----------------------------------------------------------------------------------------------+
|setparams 'OS10-A'                                                                             |
|                                                                                               |
| set root='(hd0,gpt7)'                                                                         |
| echo    'Loading OS10 ...'                                                                    |
| linux   (hd0,gpt7)/boot/os10.linux console=ttyS0,115200 root=/dev/sda7 \rw init=/bin/bash     |
| initrd  (hd0,gpt7)/boot/os10.initrd                                                           |
+-----------------------------------------------------------------------------------------------+
```

6. Press **Ctrl + x** to reboot your system. If **Ctrl + x** does not cause the system to reboot, press **Alt + 0**. The system boots up to a root shell without a password.
7. Enter `usermod -s /bin/bash linuxadmin` at the system prompt to enable `linuxadmin` user.

```
root@OS10: /# usermod -s /bin/bash linuxadmin
```

8. Verify `linuxadmin` user's password status and unlock, if locked (indicated by L, in the second column of passwd -S).

```
root@OS10:~# passwd -S linuxadmin
linuxadmin L 10/01/2018 0 99999 7 -1

root@OS10:~# passwd -u linuxadmin
passwd: password expiry information changed.
```

9. Enter `linuxadmin` for the username at the system prompt.

```
root@OS10: /# linuxadmin
```

10. Enter your password at the system prompt, then enter the new password twice.

```
root@OS10: /# passwd linuxadmin
Enter new UNIX password: xxxxxxxxx
Retype new UNIX password: xxxxxxxxxx
passwd: password updated successfully
```

11. Enter the `sync` command to save the new password.

```
root@OS10: /# sync
```

12. Reboot the system, then enter your new password.

```
root@OS10:~# reboot -f
Rebooting.
[ 3466.946967] reboot: Restarting system
```

```
BIOS Boot Selector for S5148F
Primary BIOS Version 3.36.0.1-2

SMF Version: MSS 1.2.2, FPGA 0.1
Last POR=0x11, Reset Cause=0x55
```

Restore factory defaults

To restore your system factory defaults, reboot the system to `ONIE: Uninstall OS` mode.

CAUTION: Restoring factory defaults erases any installed operating system and requires a long time to erase storage.

If it is not possible to restore your factory defaults with the installed OS, reboot the system from the Grub menu and select `ONIE: Rescue`. ONIE Rescue bypasses the installed operating system and boots the system into ONIE until you reboot the system. After ONIE Rescue completes, the system resets and boots to the ONIE console.

1 Restore the factory defaults on your system from the Grub menu using the `ONIE: Uninstall OS` command. To select which entry is highlighted, use the up and down arrow keys.

```
 +--------------------------------------+
 | ONIE: Install OS                     |
 | ONIE: Rescue                         |
 |*ONIE: Uninstall OS                   |
 | ONIE: Update ONIE                    |
 | ONIE: Embed ONIE                     |
 | ONIE: Diag ONIE                      |
 +--------------------------------------+
```

2 Press **Enter** to activate the console.

3 Return to the default ONIE settings using the `onie-uninstaller` command.

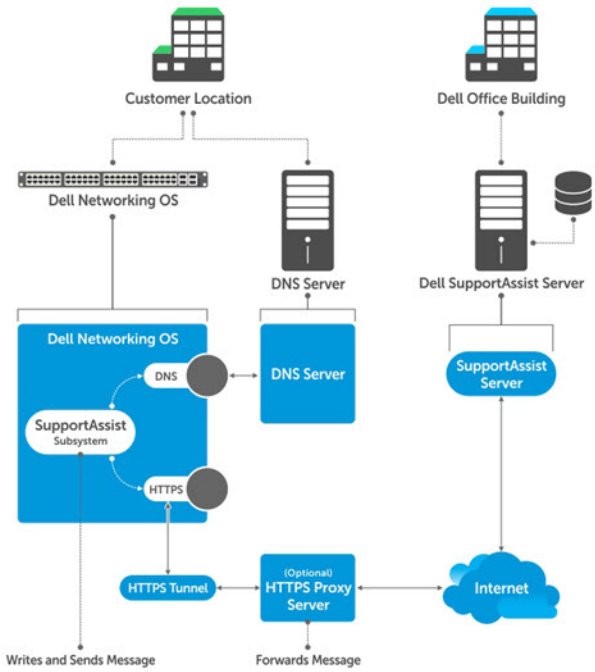
```
ONIE:/ # onie-uninstaller
uninstallerErasing internal mass storage device: /dev/sda4 (32MB)
   Percent complete: 100%
Erase complete.
Deleting partition 4 from /dev/sda
Erasing internal mass storage device: /dev/sda5 (300MB)
   Percent complete: 100%
Erase complete.
Deleting partition 5 from /dev/sda
Erasing internal mass storage device: /dev/sda6 (300MB)
   Percent complete: 100%
Erase complete.
Deleting partition 6 from /dev/sda
Erasing internal mass storage device: /dev/sda7 (12461MB)
   Percent complete: 100%
Erase complete.
Deleting partition 7 from /dev/sda
Installing for i386-pc platform.
Installation finished. No error reported.
Uninstall complete. Rebooting...
ONIE:/ # discover: Rescue mode detected. No discover stopped.
Stopping: dropbear ssh daemon... done.
Stopping: telnetd... done.
Stopping: syslogd... done.
Info: Unmounting kernel filesystems
The system is going down NOW!
Sent SIGTERM to all processes
Sent SIGKILL tosd 4:0:0:0: [sda] Synchronizing SCSI cache
Restarting system.
machine restart
```

SupportAssist

By default, SupportAssist is enabled. SupportAssist sends troubleshooting data securely to Dell EMC Technical Support. SupportAssist does not support automated email notification at the time of hardware fault alert, automatic case creation, automatic part dispatch, or reports.

To disable SupportAssist, use the `eula-consent support-assist reject` command.

SupportAssist Process



Configure SupportAssist

SupportAssist starts by default. If you do not accept end user license agreement (EULA), SupportAssist is disabled.

1 Enter SupportAssist mode from CONFIGURATION mode.

```
support-assist
```

2 (Optional) Configure the SupportAssist server URL or IP address in SUPPORT-ASSIST mode.

```
server url server-url
```

3 (Optional) Configure the interface used to connect to the SupportAssist server in SUPPORT-ASSIST mode.

```
source-interface interface
```

4 (Optional) Configure the contact information for your company in SUPPORT-ASSIST mode.

```
contact-company name {company-name}
```

5 (Optional) Configure a proxy to reach the SupportAssist server in SUPPORT-ASSIST mode.

```
proxy-server ip {ipv4-address | ipv6-address} port port-number [username user-name password password]
```

6 Trigger an activity immediately or at a scheduled time in SUPPORT-ASSIST mode.

```
do support-assist activity full-transfer {start—now | schedule [hourly | daily | weekly | monthly | yearly]}
```

Configure SupportAssist

```
OS10(config)# support-assist
OS10(conf-support-assist)# contact-company name Eureka
OS10(conf-support-assist-Eureka)# exit
OS10(conf-support-assist)# server url http://eureka.com:701
OS10(conf-support-assist)# do support-assist-activity full-transfer start-now
```

Remove SupportAssist schedule

```
OS10# no support-assist activity full-transfer schedule
```

Show EULA license

```
OS10# show support-assist eula
I accept the terms of the license agreement. You can reject the license agreement by
configuring this command 'eula-consent support-assist reject.'
By installing SupportAssist, you allow Dell to save your contact information (e.g. name, phone
number and/or email address) which would be used to provide technical support for your Dell
products and services.  Dell may use the information for providing recommendations to improve
your IT infrastructure.
Dell SupportAssist also collects and stores machine diagnostic information, which may include
but is not limited to configuration information, user supplied contact information, names of
data volumes, IP addresses, access control lists, diagnostics & performance information,
network configuration information, host/server configuration & performance information and
related data ("Collected Data") and transmits this information to Dell.  By downloading
SupportAssist and agreeing to be bound by these terms and the Dell end user license agreement,
available at: www.dell.com/aeula, you agree to allow Dell to provide remote monitoring services
of your IT environment and you give Dell the right to collect the Collected Data in accordance
with Dell's Privacy Policy, available at: www.dell.com/privacypolicycountryspecific, in order
to enable the performance of all of the various functions of SupportAssist during your
entitlement to receive related repair services from Dell. You further agree to allow Dell to
transmit and store the Collected Data from SupportAssist in accordance with these terms.  You
agree that the provision of SupportAssist may involve international transfers of data from you
to Dell and/or to Dell's affiliates, subcontractors or business partners.  When making such
transfers, Dell shall ensure appropriate protection is in plac/opt/dell/ose to safeguard the
Collected Data being transferred in connection with SupportAssist.  If you are downloading
SupportAssist on behalf of a company or other legal entity, you are further certifying to Dell
that you have appropriate authority to provide this consent on behalf of that entity.  If you
do not consent to the collection, transmission and/or use of the Collected Data, you may not
download, install or otherwise use SupportAssist.

(END)
```

Set company name

You can optionally configure name, address and territory information. Although this information is optional, it is used by Dell EMC Technical Support to identify which company owns the device.

1 (Optional) Configure contact information in SUPPORT-ASSIST mode.

```
contact-company name name
```

2 (Optional) Configure address information in SUPPORT-ASSIST mode. Use the `no address` command to remove the configuration.

```
address city name state name country name zipcode number
```

3 (Optional) Configure street address information in SUPPORT-ASSIST mode. Use double quotes to add spaces within an address. Use the `no street-address` command to remove the configuration.

```
street-address {address-line-1} [{address-line-2} {address-line-3}]
```

4 (Optional) Configure the territory and set the coverage in SUPPORT-ASSIST mode. Use the `no territory` command to remove the configuration.

```
territory company-territory
```

Configure SupportAssist company

```
OS10(conf-support-assist)# contact-company name Eureka
OS10(conf-support-assist-Eureka)# address city San Jose state California Country America
zipcode 95123
OS10(conf-support-assist-Eureka)# street-address "123 Main Street" "Bldg 999"
OS10(conf-support-assist-Eureka)# territory sales
```

Set contact information

Configure contact details in SupportAssist Company mode. You can set the name, email addresses, phone, method, and time zone. SupportAssist `contact-person` configurations are optional for the SupportAssist service.

1 (Optional) Enter the contact name in SUPPORT-ASSIST mode.

```
contact-person first firstname last lastname
```

2 Enter the email addresses in SUPPORT-ASSIST mode.

```
email-address email-address
```

3 Enter the preferred contact method in SUPPORT-ASSIST mode.

```
preferred-method {email | phone | no-contact}
```

4 Enter a contact phone number in SUPPORT-ASSIST mode.

```
phone primary number [alternate number
```

Configure contact details

```
OS10(config)# support-assist
OS10(conf-support-assist)# contact-company name Eureka
OS10(conf-support-assist-Eureka)# contact-person first John last Smith
OS10(conf-support-assist-Eureka)# email-address abc@dell.com
OS10(conf-support-assist-Eureka-JohnJamesSmith)# preferred-method email
OS10(conf-support-assist-Eureka)# phone primary 408-123-4567
```

Schedule activity

Configure the schedule for a full transfer of data. The default schedule is a full data transfer weekly — every Sunday at midnight (hour 0 minute 0).

- Configure full-transfer or log-transfer activities in EXEC mode.

  ```
  support-assist-activity {full—transfer} schedule {hourly | daily | weekly | monthly | yearly}
  ```

  - `hourly min number` — Enter the time to schedule an hourly task, from 0 to 59.
  - `daily hour number min number` — Enter the time to schedule a daily task, from 0 to 23 and 0 to 59.
  - `weekly day-of—week number hour number min number` — Enter the time to schedule a weekly task, from 0 to 6, 0 to 23, and 0 to 59.
  - `monthly day number hour number min number` — Enter the time to schedule a monthly task, from 1 to 31, 0 to 23, and 0 to 59.
  - `yearly month number day number hour number min number` — Enter the time to schedule a yearly task, from 1 to 12, 1 to 31, 0 to 23, and 0 to 59.

Configure activity schedule for full transfer

```
OS10# support-assist-activity full-transfer schedule daily hour 22 min 50
OS10# support-assist-activity full-transfer schedule weekly day-of-week 6 hour 22 min 30
OS10# support-assist-activity full-transfer schedule monthly day 15 hour 12 min 30
OS10# support-assist-activity full-transfer schedule yearly month 6 day 12 hour 6 min 30
```

Set default activity schedule

```
OS10(conf-support-assist)# no support-assist-activity full-transfer schedule
```

View status

View the SupportAssist configuration status, details, and EULA information using the `show` commands.

1 View the SupportAssist activity in EXEC mode.

```
show support-assist status
```

2 View the EULA license agreement in EXEC mode.

```
show support-assist eula
```

View SupportAssist status

```
OS10# show support-assist status
EULA            : Accepted
Service         : Enabled
Contact-Company : DellCMLCAEOS10
Street Address  : 7625 Smetana Lane Dr
                  Bldg 7615
                  Cube F577
City            : Minneapolis
State           : Minnesota
Country         : USA
Zipcode         : 55418
Territory       : USA
Contact-person  : Michael Dale
Email           : abc@dell.com
Primary phone   : 555-123-4567
Alternate phone :
Contact method  : email
Server(configured) : https://web.dell.com
Proxy IP        :
Proxy Port      :
Proxy username  :
Activity Enable State   :
     Activity           State
  -------------------------------
    coredump-transfer   enabled
    event-notification  enabled
    full-transfer       enabled

Scheduled Activity List   :
 Activity       Schedule                Schedule created on
 ------------------------------------------------------------
 full-transfer  weekly: on sun at 00:00  Sep 12,2016 18:57:40

Activity Status   :
   Activity           Status          last start         last success
  ------------------------------------------------------------------------
   coredump-transfer  success    Sep 12,2016 20:48:41    Sep 12,2016 20:48:42
   event-notification success    Sep 12,2016 20:51:51    Sep 12,2016 20:51:51
   full-transfer      success    Sep 12,2016 20:30:28    Sep 12,2016 20:30:52
```

View EULA license

```
OS10# show support-assist eula
I accept the terms of the license agreement. You can reject the license agreement by
configuring this command 'eula-consent support-assist reject.'
By installing SupportAssist, you allow Dell to save your contact information (e.g. name, phone
number and/or email address) which would be used to provide technical support for your Dell
products and services.  Dell may use the information for providing recommendations to improve
your IT infrastructure.
Dell SupportAssist also collects and stores machine diagnostic information, which may include
```

```
but is not limited to configuration information, user supplied contact information, names of
data volumes, IP addresses, access control lists, diagnostics & performance information,
network configuration information, host/server configuration & performance information and
related data ("Collected Data") and transmits this information to Dell.  By downloading
SupportAssist and agreeing to be bound by these terms and the Dell end user license agreement,
available at: www.dell.com/aeula, you agree to allow Dell to provide remote monitoring services
of your IT environment and you give Dell the right to collect the Collected Data in accordance
with Dell's Privacy Policy, available at: www.dell.com/privacypolicycountryspecific, in order
to enable the performance of all of the various functions of SupportAssist during your
entitlement to receive related repair services from Dell,.  You further agree to allow Dell to
transmit and store the Collected Data from SupportAssist in accordance with these terms.  You
agree that the provision of SupportAssist may involve international transfers of data from you
to Dell and/or to Dell's affiliates, subcontractors or business partners.  When making such
transfers, Dell shall ensure appropriate protection is in place to safeguard the Collected Data
being transferred in connection with SupportAssist.  If you are downloading SupportAssist on
behalf of a company or other legal entity, you are further certifying to Dell that you have
appropriate authority to provide this consent on behalf of that entity.  If you do not consent
to the collection, transmission and/or use of the Collected Data, you may not download, install
or otherwise use SupportAssist.

(END)
```

SupportAssist commands

activity

Enables SupportAssist activities to run on an associated trigger or schedule time.

Syntax `activity [coredump-transfer | event-notification | full-transfer} enable`

Parameters

- `coredump-transfer` — Enables transfer of core dump files.
- `event-notification` — Enables transfer of event notification files.
- `full-transfer` — Enables transfer of logs and technical support information.

Default Enabled

Command Mode SUPPORT-ASSIST

Usage Information Use the `no` version of this command to remove the configuration.

Example (Event)
```
OS10(conf-support-assist)# activity event-notification enable
```

Example (Full)
```
OS10(conf-support-assist)# activity full-transfer enable
```

Example (Turn Off)
```
OS10(conf-support-assist)# no activity coredump-transfer enable
```

Supported Releases 10.2.0E or later

contact-company

Configures the company contact information.

Syntax `contact-company` *`name`*

Parameters *`name`* — Enter the contact company name. A maximum of 140 characters.

Default	Not configured
Command Mode	SUPPORT-ASSIST
Usage Information	You can enter only one contact-company. Use double quotes to enclose additional contact information. The `no` version of this command removes the configuration.
Example	`OS10(conf-support-assist)# contact-company name Eureka` `OS10(conf-support-assist-Eureka)#`
Supported Releases	10.2.0E or later

contact-person

Configures the contact name for an individual.

Syntax	`contact-person [first` *`firstname`* `last` *`lastname`*`]`
Parameters	• `first` *`firstname`* — Enter the keyword and the first name for the contact person. Use double quotes for more than one first name. • `last` *`lastname`* — Enter the keyword and the last name for the contact person.
Default	Not configured
Command Mode	SUPPORT-ASSIST
Usage Information	The `no` version of this command removes the configuration.
Example	`OS10(conf-support-assist-Eureka)# contact-person first "John James" last Smith`
Supported Releases	10.2.0E or later

email-address

Configures the email address for the contact name.

Syntax	`email—address` *`address`*
Parameters	*`address`* — Enter the email address for the contact name.
Default	Not configured
Command Mode	SUPPORT-ASSIST
Usage Information	The `no` version of this command removes the configuration.
Example	`OS10(conf-support-assist-Eureka-JohnJamesSmith)# email-address jjsmith@eureka.com`
Supported Releases	10.2.0E or later

eula-consent

Accepts or rejects the SupportAssist end-user license agreement (EULA).

Syntax	`eula—consent {support-assist} {accept | reject}`

Parameters

- `support-assist` — Enter to accept or reject the EULA for the service.
- `accept` — Enter to accept the EULA-consent.
- `reject` — Enter to reject EULA-consent.

Default Not configured

Command Mode CONFIGURATION

Usage Information If you reject the end-user license agreement, you cannot access Configuration mode. If there is an existing SupportAssist configuration, the configuration is not removed and the feature is disabled.

Example (Accept)

```
OS10(config)# eula-consent support-assist accept
```

Example (Reject)

```
OS10(config)# eula-consent support-assist reject

This action will disable Support Assist and erase all configured data.Do you
want to proceed ? [Y/N]:Y
```

Supported Releases 10.2.0E or later

preferred-method

Configures a preferred method to contact an individual.

Syntax `preferred-method {email | phone | no-contact}`

Parameters

- `email` — Enter to select email as the preferred contact method.
- `phone` — Enter to select phone as the preferred contact method.
- `no-contact` — Enter to select no-contact as the preferred contact method.

Default No-contact

Command Mode SUPPORT-ASSIST

Usage Information The `no` version of this command removes the configuration.

Example

```
OS10(conf-support-assist-Eureka-JohnJamesSmith)# preferred-method email
```

Supported Releases 10.2.0E or later

proxy-server

Configures a proxy IP address for reaching the SupportAssist server.

Syntax `proxy-server ip` *`ipv4-address`* `port` *`number`*

Parameters

- *`ipv4-address`*— Enter the IPv4 address of the proxy server in a dotted decimal format (A.B.C.D).
- *`number`* — Enter the port number, from 0 to 65535.

Default Not configured

Command Mode SUPPORT-ASSIST

Usage Information You cannot use an IPv6 address with this command.

Example

```
OS10(conf-support-assist)# proxy-server ip 10.1.1.5 port 701
```

Supported Releases 10.2.0E or later

server url

Configures the domain or IP address of the remote SupportAssist server.

Syntax `server url server-url-string`

Parameters *server-url-string* — Enter the domain or IP address of the remote SupportAssist server. To include a space, enter a space within double quotes.

Default https://stor.g3.ph.dell.com

Command Mode SUPPORT-ASSIST

Usage Information Only configure one SupportAssist server. If you do not configure the SupportAssist server, the system uses the non-configurable default server. Use the `show support-assist status` command to view the server configuration. The `no` version of this command removes the remote server.

Example

```
OS10(conf-support-assist)# server url https://eureka.com:444
```

Supported Releases 10.2.0E or later

show support-assist eula

Displays the EULA for SupportAssist.

Syntax `show support-assist eula`

Parameters None

Default None

Command Mode EXEC

Usage Information Use the `eula-consent support-assist accept` command to accept the license agreement.

Example

```
OS10# show support-assist eula
I accept the terms of the license agreement. You can reject the license
agreement by configuring this command 'eula-consent support-assist reject.'
By installing SupportAssist, you allow Dell, Inc. to save your contact
information (e.g. name, phone number and/or email address) which would be used
to provide technical support for your Dell, Inc. products and services.  Dell,
Inc. may use the information for providing recommendations to improve your IT
infrastructure.
SupportAssist also collects and stores machine diagnostic information, which
may include but is not limited to configuration information, user supplied
contact information, names of data volumes, IP addresses, access control
lists, diagnostics & performance information, network configuration
information, host/server configuration & performance information and related
data ("Collected Data") and transmits this information to Dell, Inc.  By
downloading SupportAssist and agreeing to be bound by these terms and the
Dell, Inc. end user license agreement, available at: www.dell.com/aeula, you
agree to allow Dell, Inc. to provide remote monitoring services of your IT
environment and you give Dell, Inc. the right to collect the Collected Data in
accordance with Dell, Inc.'s Privacy Policy, available at: www.dell.com/
privacypolicycountryspecific, in order to enable the performance of all of the
various functions of SupportAssist during your entitlement to receive related
repair services from Dell, Inc.  You further agree to allow Dell, Inc. to
transmit and store the Collected Data from SupportAssist in accordance with
```

```
these terms.  You agree that the provision of SupportAssist may involve
international transfers of data from you to Dell, Inc. and/or to Dell, Inc.'s
affiliates, subcontractors or business partners.  When making such transfers,
Dell, Inc. shall ensure appropriate protection is in place to safeguard the
Collected Data being transferred in connection with SupportAssist. If you are
downloading SupportAssist on behalf of a company or other legal entity, you
are further certifying to Dell, Inc. that you have appropriate authority to
provide this consent on behalf of that entity.  If you do not consent to the
collection, transmission and/or use of the Collected Data, you may not
download, install or otherwise use SupportAssist.

(END)
```

Supported Releases 10.2.0E or later

show support-assist status

Displays SupportAssist status information including activities and events.

Syntax `show support-assist status`

Parameters None

Default Not configured

Command Mode EXEC

Example

```
OS10# show support-assist status
EULA             : Accepted
Service          : Enabled
Contact-Company  : DellCMLCAEOS10
Street Address   : 7625 Smetana Lane Dr
                   Bldg 7615
                   Cube F577
City             : Minneapolis
State            : Minnesota
Country          : USA
Zipcode          : 55418
Territory        : USA
Contact-person   : Michael Dale
Email            : abc@dell.com
Primary phone    : 555-123-4567
Alternate phone  :
Contact method   : email
Server(configured) : https://web.dell.com
Proxy IP         :
Proxy Port       :
Proxy username   :
Activity Enable State   :
     Activity           State
  ---------------------------------
    coredump-transfer   enabled
    event-notification  enabled
    full-transfer       enabled

Scheduled Activity List   :
  Activity        Schedule                Schedule created on
  ------------------------------------------------------------
 full-transfer   weekly: on sun at 00:00   Sep 12,2016 18:57:40

Activity Status   :
  Activity           Status    last start              last success
  ------------------------------------------------------------------------
  coredump-transfer  success   Sep 12,2016 20:48:41     Sep 12,2016 20:48:42
```

```
event-notification success   Sep 12,2016 20:51:51     Sep 12,2016 20:51:51
full-transfer      success   Sep 12,2016 20:30:28     Sep 12,2016 20:30:52
```

Supported Releases 10.2.0E or later

source-interface

Configures the interface used to connect to the SupportAssist server.

Syntax `source-interface interface`

Parameters *interface*:

- `ethernet node/slot/port[:subport]` — Enter a physical Ethernet interface.
- `loopback number` — Enter a Loopback interface, from 0 to 16383.
- `management 1/1/1` — Enter the management interface.
- `port-channel channel-id` — Enter a port-channel ID, from 1 to 28.
- `vlan vlan-id` — Enter a VLAN ID, from 1 to 4093.

Default A source interface is not configured.

Command Mode SUPPORT-ASSIST

Usage Information The `no` version of this command removes the configured source interface.

Example

```
OS10(conf-support-assist)# source-interface ethernet 1/1/4
```

Supported Releases 10.4.0E(R1) or later

street-address

Configures the street address information for the company.

Syntax `street-address {address}`

Parameters *address* — Enter one or more addresses in double quotes. A maximum of 140 characters.

Default Not configured

Command Mode SUPPORT-ASSIST

Usage Information Add spaces to the company street address by enclosing the address in quotes. Separate each address with a space to place on a new line. The `no` version of this command removes the company address configuration.

Example

```
OS10(conf-support-assist-Eureka)# street-address "One Dell Way" "Suite 100"
```

Supported Releases 10.2.0E or later

support-assist-activity

Schedules a time to transfer the activity log.

Syntax `support-assist-activity full-transfer [start-now] [schedule {hourly minute | daily hour number min number | weekly day-of-week number hour number | monthly day number hour number min number | yearly month number day number}]`

Parameters

- `start-now` — Schedules the transfer to start immediately.
- `hourly` *`minute`* — Schedule an hourly task, from 0 to 59.
- `daily` — Schedule a daily task:
  - `hour` *`number`* — Enter the keyword and number of hours to schedule the daily task, from 0 to 23.
  - `min` *`number`* — Enter the keyword and number of minutes to schedule the daily task, from 0 to 59.
- `weekly` — Schedule a weekly task:
  - `day-of-week` *`number`* — Enter the keyword and number for the day of the week to schedule the task, from 0 to 6.
  - `hour` *`number`* — Enter the keyword and number of the hour to schedule the weekly task, from 0 to 23.
- `monthly` — Schedule a monthly task:
  - `day` *`number`* — Enter the number for the day of the month to schedule the task, from 1 to 31.
  - `hour` *`number`* — Enter the number for the hour of the day to schedule the task, from 0 to 23.
  - `min` *`number`* — Enter the number for the minute of the hour to schedule the task, from 0 to 59.
- `yearly` — Schedule the yearly task:
  - `month` *`number`* — Enter the keyword and number of the month to schedule the yearly task, from 1 to 12.
  - `day` *`number`* — Enter the keyword and the number of the day to schedule the monthly task, from 1 to 31.

Default Weekly on Sunday at midnight (hour 0 minute 0)

Command Mode EXEC

Usage Information The `no` version of this command removes the schedule activity.

Example

```
OS10# support-assist-activity full-transfer schedule daily hour 22 min 50
```

Supported Releases 10.2.0E or later

territory

Configures the territory for the company.

Syntax `territory` *`territory`*

Parameters *`territory`* — Enter the territory for the company.

Default Not configured

Command Mode CONFIG-SUPPORT-ASSIST

Usage Information The `no` version of this command removes the company territory configuration.

Example

```
OS10(conf-support-assist)# contact-company name Eureka
OS10(conf-support-assist-Eureka)# territory west
```

Supported Releases 10.2.0E or later

Support bundle

The Support Bundle is based on the `sosreport` tool. Use the Support Bundle to generate an `sosreport` tar file that collects Linux system configuration and diagnostics information, as well as the `show` command output to send to Dell EMC Technical Support.

To send Dell EMC Technical Support troubleshooting details about the Linux system configuration and OS10 diagnostics, generate an `sosreport` tar file.

1 Generate the tar file in EXEC mode.
```
generate support-bundle
```
2 Verify the generated file in EXEC mode.
```
dir supportbundle
```
3 Send the support bundle using FTP/SFTP/SCP/TFTP in EXEC mode.
```
copy supportbundle://sosreport-filename.tar.gz tftp://server-address/path
```

Use the `delete supportbundle://sosreport-filename.tar.gz` command to delete a generated support bundle.

Event notifications

Event notifications for the `generate support-bundle` command process at the start and end of the bundle they support, and reports either success or failure.

Support bundle generation start event

```
Apr 19 16:57:55: %Node.1-Unit.1:PRI:OS10 %log-notice:SUPPORT_BUNDLE_STARTED: generate support-
bundle execution has started successfully:All Plugin options disabled
Apr 19 16:57:55: %Node.1-Unit.1:PRI:OS10 %log-notice:SUPPORT_BUNDLE_STARTED: generate support-
bundle execution has started successfully:All Plugin options enabled
```

sosreport generation start event

```
May 11 22:9:43: %Node.1-Unit.1:PRI:OS10 %log-notice:SOSREPORT_GEN_STARTED: CLI output
collection task completed; sosreport execution task started:All Plugin options disabled
May 11 22:9:43: %Node.1-Unit.1:PRI:OS10 %log-notice:SOSREPORT_GEN_STARTED: CLI output
collection task completed; sosreport execution task started:All Plugin options enabled
```

Support bundle generation successful event

```
Apr 19 17:0:9: %Node.1-Unit.1:PRI:OS10 %log-notice:SUPPORT_BUNDLE_COMPLETED: generate support-
bundle execution has completed successfully:All Plugin options disabled
Apr 19 17:0:9: %Node.1-Unit.1:PRI:OS10 %log-notice:SUPPORT_BUNDLE_COMPLETED: generate support-
bundle execution has completed successfully:All Plugin options enabled
```

Support bundle generation failure

```
Apr 19 17:0:14: %Node.1-Unit.1:PRI:OS10 %log-notice:SUPPORT_BUNDLE_FAILURE: Failure in generate
support-bundle execution:All Plugin options disabled
Apr 19 17:0:14: %Node.1-Unit.1:PRI:OS10 %log-notice:SUPPORT_BUNDLE_FAILURE: Failure in generate
support-bundle execution:All Plugin options enabled
```

generate support-bundle

Generates an `sosreport` tar file that collects configuration and diagnostic information on Linux systems.

Syntax	`generate support-bundle [enable-all-plugin-options]`
Parameters	`enable-all-plugin-options` — (Optional) Generate a full support bundle with all plugin options enabled.
Defaults	None
Command Mode	EXEC

Usage Information	To send the tar file to Dell EMC Technical Support, use the `dir supportbundle` and `copy supportbundle://sosreport-OS10-file-number.tar.gz tftp://server-address/path` commands.
Example	`OS10# generate support-bundle`
Example (Enable Options)	`OS10# generate support-bundle enable-all-plugin-options`
Supported Releases	10.2.0E or later

System monitoring

Monitor OS10 using system alarms and log information.

System alarms

Alarms alert you to conditions that might prevent normal device operation:

- **Critical** — A critical condition exists and requires immediate action. A critical alarm may trigger if one or more hardware components fail, or one or more hardware components exceeds temperature thresholds.
- **Major** — A major error occurred and requires escalation or notification. For example, a major alarm may trigger if an interface failure occurs, such as a port-channel being down.
- **Minor** — A minor error or non-critical condition occurred that, if left unchecked, might cause system service interruption or performance degradation. A minor alarm requires monitoring or maintenance.
- **Informational** — An informational error occurred but does not impact performance. Monitor an informational alarm until the condition changes.

Triggered alarms are in one of these states:

- **Active** — Alarms that are current and not cleared.
- **Cleared** — Alarms that are resolved and the device has returned to normal operation.

System logging

You can change the system logging default settings using the severity level to control the type of system messages that log. The range of logging severities are:

- `log-emerg` — System is unstable.
- `log-alert` — Immediate action needed.
- `log-crit` — Critical conditions.
- `log-err` — Error conditions.
- `log-warning` — Warning conditions.
- `log-notice` — Normal but significant conditions, the default.
- `log-info` — Informational messages.
- `log-debug` — Debug messages.

- Enter the minimum severity level for logging to the console in CONFIGURATION mode.
  `logging console severity`
- Enter the minimum severity level for logging to the system log file in CONFIGURATION mode.
  `logging log-file severity`

- Enter the minimum severity level for logging to terminal lines in CONFIGURATION mode.
  ```
  logging monitor severity
  ```
- Enter which server to use for syslog messages with the hostname or IP address in CONFIGURATION mode.
  ```
  logging server {hostname/ip-address severity}
  ```

Disable system logging

You can use the `no` version of any logging command to disable system logging.

- Disable console logging and reset the minimum logging severity to the default in CONFIGURATION mode.
  ```
  no logging console severity
  ```
- Disable log-file logging and reset the minimum logging severity to the default in CONFIGURATION mode.
  ```
  no logging log-file severity
  ```
- Disable monitor logging and reset the minimum logging severity to the default in CONFIGURATION mode.
  ```
  no logging monitor severity
  ```
- Disable server logging and reset the minimum logging severity to the default in CONFIGURATION mode.
  ```
  no logging server severity
  ```
- Re-enable any logging command in CONFIGURATION mode.
  ```
  no logging enable
  ```

Enable server logging for log notice

```
OS10(config)# logging server dell.com severity log-notice
```

View system logs

The system log-file contains system event and alarm logs.

Use the `show trace` command to view the current syslog file. All event and alarm information is sent to the syslog server, if one is configured.

The `show logging` command accepts the following parameters:

- `log-file` — Provides a detailed log including both software and hardware saved to a file.
- `process-names` — Provides a list of all processes currently running which can be filtered based on the process-name.

View logging log-file

```
OS10# show logging log-file
Jun  1 05:01:46 %Node.1-Unit.1:PRI:OS10 %log-notice:ETL_SERVICE_UP: ETL service
is up
Jun  1 05:02:06 %Node.1-Unit.1:PRI:OS10 %log-notice:EQM_UNIT_DETECTED: Unit pres
ent:Unit 1#003
Jun  1 05:02:09 %Node.1-Unit.1:PRI:OS10 %log-notice:EQM_PSU_DETECTED: Power Supp
ly Unit present:PSU 1#003
Jun  1 05:02:09 %Node.1-Unit.1:PRI:OS10 %log-notice:EQM_PSU_DETECTED: Power Supp
ly Unit present:PSU 2#003
Jun  1 05:02:09 %Node.1-Unit.1:PRI:OS10 %log-notice:EQM_FAN_TRAY_DETECTED: Fan t
ray present:Fan tray 1#003
Jun  1 05:02:09 %Node.1-Unit.1:PRI:OS10 %log-notice:EQM_FAN_TRAY_DETECTED: Fan t
ray present:Fan tray 2#003
Jun  1 05:02:09 %Node.1-Unit.1:PRI:OS10 %log-notice:EQM_FAN_TRAY_DETECTED: Fan t
ray present:Fan tray 3#003
Jun  1 05:02:09 %Node.1-Unit.1:PRI:OS10 %log-crit:EQM_FAN_AIRFLOW_MISMATCH: MAJO
R ALARM: FAN AIRFLOW MISMATCH: SET: One or more fans have mismatching or unknown
 airflow directions#003
Jun  1 05:02:10 %Node.1-Unit.1:PRI:OS10 %log-notice:NDM_SERVICE_UP: NDM Service
Ready!
```

```
Jun  1 05:02:10 %Node.1-Unit.1:PRI:OS10 %log-notice:SU_SERVICE_UP: Software upgr
ade service is up:software upgrade service up
--More--
```

View logging process names

```
OS10# show logging process-names
dn_alm
dn_app_vlt
dn_app_vrrp
dn_bgp
dn_dot1x
dn_eqa
dn_eqm
dn_eth_drv
dn_etl
dn_i3
dn_ifm
dn_infra_afs
dn_issu
dn_l2_services
dn_l2_services_
dn_l2_services_
dn_l2_services_
dn_l2_services_
dn_l3_core_serv
dn_l3_service
dn_lacp
dn_lldp
dn_mgmt_entity_
--More--
```

Environmental monitoring

Monitors the hardware environment to detect temperature, CPU, and memory utilization.

View environment

```
OS10# show environment

Unit    State             Temperature Voltage
-------------------------------------------
1       up                42
-------------------------------------------
Thermal sensors
Unit   Sensor-Id        Sensor-name          Temperature
-----------------------------------------------------
1       1           T2 temp sensor              28
1       2           system-NIC temp sensor      25
1       3           Ambient temp sensor         24
1       4           NPU temp sensor             40
-----------------------------------------------------
```

Link-bundle monitoring

Monitoring link aggregation group (LAG) bundles allows the traffic distribution amounts in a link to look for unfair distribution at any given time. A threshold of 60% is an acceptable amount of traffic on a member link.

Links are monitored in 15-second intervals for three consecutive instances. Any deviation within that time sends syslog and an alarm event generates. When the deviation clears, another syslog sends and a clear alarm event generates.

Link-bundle utilization calculates the total bandwidth of all links divided by the total bytes-per-second of all links. If you enable monitoring, the utilization calculation performs when the utilization of the link-bundle (not a link within a bundle) exceeds 60%.

Configure Threshold level for link-bundle monitoring

```
OS10(config)# link-bundle-trigger-threshold 10
```

View link-bundle monitoring threshold configuration

```
OS10(config)# do show running-configuration
link-bundle-trigger-threshold 10
!
...
```

Show link-bundle utilization

```
OS10(config)# do show link-bundle-utilization

Link-bundle trigger threshold - 10
```

Alarm commands

alarm acknowledge

Acknowledges an active alarm.

Syntax	`alarm acknowledge` *`sequence-number`*
Parameters	• *`sequence-number`* — Acknowledge the alarm corresponding to the sequence number.
Default	Not configured
Command Mode	EXEC
Usage Information	Use the `show alarm` command to view all active alarms. Use active alarm sequence numbers to acknowledge specific alarms.
Example	`OS10# alarm acknowledge 1`
Supported Releases	10.4.3 or later

alarm clear

Clears the alarm based on the alarm index for a user-clearable alarm. a transient alarm.

Syntax	`alarm clear` *`alarm-index`*
Parameters	`clear` *`alarm-index`* — Enter the alarm ID to clear the alarm.
Default	Not configured
Command Mode	EXEC
Usage Information	Use the `show alarm index` command to view a list of alarm IDs.
Example	`OS10# alarm clear 200`

Supported Releases 10.2.0E or later

show alarms

Displays all current active alarms in the system.

Syntax `show alarms`

Parameters None

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show alarms
Sq No  Severity   Name                Timestamp                 Source
-----  --------   ------------------  ------------------------  ------
  3    major      EQM_MORE_PSU_FAULT  Sun 10-07-2018 18:39:47  /psu/2
* 4    Minor      EQM_MORE_PSU_FAULT  Sun 10-07-2018 18:39:47  /psu/1
```

Supported Releases 10.2.0E or later

show alarms details

Displays details about active alarms.

Syntax `show alarms details`

Parameters None

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show alarms details
Active-alarm details - 1
-----------------------------------------
Sequence Number:   3
Severity:          major
Source:            /psu/2
Name:              EQM_MORE_PSU_FAULT
Description:       psu 2 is not working correctly
Raise-time:        Sun 10-07-2018 18:39:47
Ack-time:
State:             raised
-----------------------------------------

Active-alarm details - 2
-----------------------------------------
Sequence Number:   4
Severity:          minor
Source:            /psu/1
Name:              EQM_MORE_PSU_FAULT
Description:       psu 2 is not working correctly
Raise-time:        Sun 10-07-2018 18:39:47
Ack-time:          Sun 10-07-2018 20:39:47
New:               true
State:             acknowledged
-----------------------------------------
```

Supported Releases 10.2.0E or later

show alarms history

Displays the history of all alarm events.

Syntax `show alarms history [summary]`

Parameters `summary` — Enter to view a summary of the alarm history.

Default None

Command Mode EXEC

Usage Information None

Example

```
OS10# show alarms history
Sq No    State      Name                Timestamp                 Source
-----    --------   ------------------  ------------------------  ------
  1      Stateless  SYSTEM_REBOOT       Sun 10-07-2018 15:39:47
  2      Raised     EQM_FANTRAY_FAULT   Sun 10-07-2018 16:39:47  /fantray/3
  3      Raised     EQM_MORE_PSU_FAULT  Sun 10-07-2018 18:39:47  /psu/2
  4      Raised     EQM_MORE_PSU_FAULT  Sun 10-07-2018 18:39:47  /psu/1
  5      Ack        EQM_MORE_PSU_FAULT  Sun 10-07-2018 20:39:47  /psu/1
  6      Cleared    EQM_FANTRAY_FAULT   Sun 10-07-2018 22:39:47  /fantray/3
```

Example (Summary)

```
OS10# show alarms history summary
Alarm History Summary
---------------------
Total-count:      4
Raised-count:     3
Ack-count:        1
Cleared-count:    1
Stateless-count:  1
-----------------------------------------
```

Supported Releases 10.4.3 or later

show alarms sequence

Displays information corresponding to the active alarm based on the sequence number that you specify.

Syntax `show alarms sequence` *`sequence-number`*

Parameters

- `sequence-number` — Enter the sequence number corresponding to the active alarm.

Default None

Command Mode EXEC

Usage Information Use the show alarms command to view all active alarms. Use an active alarm sequence number to view detailed information about that alarm.

Example

```
NOS# show alarms sequence 3
Active-alarm details - 1
-----------------------------------------
Sequence Number:   3
Severity:          major
Type:              1081375
Source:            /psu/2
Name:              EQM_MORE_PSU_FAULT
```

```
Description:       psu 2 is not working correctly
Raise-time:        Sun 10-07-2018 18:39:47
Ack-time:
State:             raised
-------------------------------------------
```

Supported Releases 10.4.3E or later

show alarms index

Displays information about a specific alarm using the alarm ID.

Syntax `show alarms index alarm-id`

Parameters `index alarm-id` — Enter the keyword and the alarm ID to view specific information.

Default Not configured

Command Mode EXEC

Usage Information Use the `alarm-id` to clear and view alarm details.

Example

```
OS10# show alarms index 1

Active-alarm details - 1
-------------------------------------------
Index:             1
Sequence Number:   5
Severity:          warning
Type:              1081364
Source:            Node.1-Unit.1
Name:              EQM_THERMAL_WARN_CROSSED
Description:
Raise-time:        Sep 20 0:16:52
Clear-time:
New:               true
State:             raised
```

Supported Releases 10.2.0E or later

show alarms severity

Displays all active alarms corresponding to a specific severity level.

Syntax `show alarms severity severity`

Parameters *severity* — Set the alarm severity:

- `critical` — Critical alarm severity.
- `major` — Major alarm severity.
- `minor` — Minor alarm severity.
- `warning` — Warning alarm severity.

Default Not configured

Command Mode EXEC

Usage Information None

Example (Warning)

```
OS10# show alarms severity warning
```

```
Active-alarm details - 1
-----------------------------------------
Sequence Number:   5
Severity:          warning
Type:              1081364
Source:            Node.1-Unit.1
Name:              EQM_THERMAL_WARN_CROSSED
Description:
Raise-time:        Sat 10-06-2018 0:1:5
Ack-time:          Sun 10-07-2018 20:39:47
New:               true
State:             raised
```

Example (Critical)

```
OS10# show alarms severity critical

Active-alarm details - 0
-----------------------------------------
Sequence Number:   1
Severity:          critical
Type:              1081367
Source:            Node.1-Unit.1
Name:              EQM_THERMAL_CRIT_CROSSED
Description:
Raise-time:        Sat 10-06-2018 0:1:5
Ack-time:          Sun 10-07-2018 20:39:47
New:               true
State:             raised
```

Example (Minor)

```
NOS# show alarms severity minor
Active-alarm details - 1
-----------------------------------------
Sequence Number:   4
Severity:          minor
Type:              1081375
Source:            /psu/1
Name:              EQM_MORE_PSU_FAULT
Description:       psu 2 is not working correctly
Raise-time:        Sun 10-07-2018 18:39:47
Ack-time:          Sun 10-07-2018 20:39:47
New:               true
State:             acknowledged
-----------------------------------------
```

Supported Releases 10.4.3 or later

show alarms summary

Displays the summary of all active alarms.

Syntax `show alarms summary`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# show alarms summary
Active-alarm Summary
-----------------------------------------
Total-count:      2
Critical-count:   0
Major-count:      1
```

```
Minor-count:      1
Warning-count:    0
-----------------------------------------
```

Supported Releases 10.2.0E or later

Logging commands

clear logging

Clears messages in the logging buffer.

Syntax `clear logging log-file`

Parameters None

Default Not configured

Command Mode EXEC

Usage Information None

Example

```
OS10# clear logging log-file

Proceed to clear the log file [confirm yes/no(default)]:
```

Supported Releases 10.2.0E or later

logging console

Disables, enables, or configures the minimum severity level for logging to the console.

Syntax `logging console {disable | enable | severity}`

Parameters *severity* — Set the minimum logging severity level:

- `log-emerg` — Set to unusable.
- `log-alert` — Set to immediate action is needed.
- `log-crit` — Set to critical conditions.
- `log-err` — Set to error conditions.
- `log-warning` — Set to warning conditions.
- `log-notice` — Set to normal but significant conditions, the default.
- `log-info` — Set to informational messages.
- `log-debug` — Set to debug messages.

Default Log-notice

Command Mode CONFIGURATION

Usage Information To set the severity to the default level, use the `no logging console severity` command. The default severity level is `log-notice`.

Example
```
OS10(config)# logging console disable
```

Example (Enable)
```
OS10(config)# logging console enable
```

Example (Severity)
```
OS10(config)# logging console severity log-warning
```

Supported Releases 10.2.0E or later

logging enable

Enables system logging.

Syntax `logging enable`

Parameters None

Default Enabled

Command Mode CONFIGURATION

Usage Information The `no` version of this command disables all logging.

Example
```
OS10(config)# logging enable
```

Supported Releases 10.2.0E or later

logging log-file

Disables, enables, or sets the minimum severity level for logging to the log file.

Syntax `logging log-file {disable | enable | severity}`

Parameters *severity* — Set the minimum logging severity level:

- `log-emerg` — Set the system as unusable.
- `log-alert` — Set to immediate action is needed.
- `log-crit` — Set to critical conditions.
- `log-err` — Set to error conditions.
- `log-warning` — Set to warning conditions.
- `log-notice` — Set to normal but significant conditions, the default.
- `log-info` — Set to informational messages.
- `log-debug` — Set to debug messages.

Default Log-notice

Command Mode CONFIGURATION

Usage Information To reset the log-file severity to the default level, use the `no logging log-file severity` command. The default severity level is `log-notice`.

Example
```
OS10(config)# logging log-file disable
```

Example (Enable)
```
OS10(config)# logging log-file enable
```

Example (Severity)
```
OS10(config)# logging log-file severity log-notice
```

Supported Releases 10.2.0E or later

logging monitor

Set the minimum severity level for logging to the terminal lines.

Syntax `logging monitor severity` *`severity-level`*

Parameters *`severity-level`* — Set the minimum logging severity level:

- `log-emerg` — Set the system as unusable.
- `log-alert` — Set to immediate action is needed.
- `log-crit` — Set to critical conditions.
- `log-err` — Set to error conditions.
- `log-warning` — Set to warning conditions.
- `log-notice` — Set to normal but significant conditions, the default.
- `log-info` — Set to informational messages.
- `log-debug` — Set to debug messages.

Default Log-notice

Command Mode CONFIGURATION

Usage Information To reset the monitor severity to the default level, use the `no logging monitor severity` command. The default severity level is `log-notice`.

Example
```
OS10(config)# logging monitor severity log-info
```

Supported Releases 10.2.0E or later

logging server

Configures the remote syslog server.

Syntax `logging server {`*`hostname`* `|` *`ipv4-address`* `|` *`ipv6-address`*`} [severity` *`severity-level`* `| vrf management [severity` *`severity-level`*`] [tcp` *`port-number`* `| udp` *`port-number`*`]`

Parameters

- *`hostname`* `|` *`ipv4-address`* `|` *`ipv6-address`* — (Optional) Enter either the hostname or IPv4/IPv6 address of the logging server.
- `vrf management` — (Optional) Configure the logging server for the management VRF instance.
- *`severity-level`* — (Optional) Set the logging threshold severity:
  - `log-emerg` — System as unusable.
  - `log-alert` — Immediate action is needed.
  - `log-crit` — Critical conditions.
  - `log-err` — Error conditions.
  - `log-warning` — Warning conditions.
  - `log-notice` — Normal but significant conditions, the default.

  - `log-info` — Informational messages.
  - `log-debug` — Debug messages.
- `tcp` *`port-number`* — (Optional) Send syslog messages over TCP to a specified port on a remote logging server, from 1 to 65535.
- `udp` *`port-number`* — (Optional) Send syslog messages over UDP to a specified port on a remote logging server, from 1 to 65535; default 514.

Defaults

System messages of the `log-notice` security level and lower are generated.

Syslog messages are sent over UDP to port 514 on a remote logging server.

Command Mode

CONFIGURATION

Usage Information

Starting from 10.3.0E or later, this command supports IPv6 addresses. The previous versions support only IPv4 addresses. The `no` version of this command deletes the syslog server.

Example

```
OS10(config)# logging server dell.com severity log-info
OS10(config)# logging server fda8:6c3:ce53:a890::2 tcp 1468
OS10(config)# logging server dell.com vrf management severity log-debug
```

Supported Releases

10.2.0E or later

show logging

Displays system logging messages by log file, process-names, or summary.

Syntax

```
show logging {log-file [process-name | line-numbers] | process-names}
```

Parameters

- *`process-name`* — (Optional) Enter the process-name to use as a filter in syslog messages.
- *`line-numbers`* — (Optional) Enter the number of lines to include in the logging messages, from 1 to 65535.

Default

None

Command Mode

EXEC

Usage Information

The output from this command is the */var/log/eventlog* file.

Example (Log File)

```
OS10# show logging log-file process-name dn_qos
```

Example (Process-Names)

```
OS10# show logging process-names
dn_pas_svc
dn_system_mgmt_
dn_env_tmpctl
dn_pm
dn_eth_drv
dn_etl
dn_eqa
dn_alm
dn_eqm
dn_issu
dn_swupgrade
dn_ifm
dn_ppm
dn_l2_services
dn_dot1x
dn_l3_core_serv
dn_policy
```

```
dn_qos
dn_switch_res_m
dn_ospfv3
dn_lacp
dn_i3
dn_supportassis
--More--
```

Supported Releases 10.2.0E or later

show trace

Displays trace messages.

Syntax `show trace [number-lines]`

Parameters *number-lines* — (Optional) Enter the number of lines to include in log messages, from 1 to 65535.

Default Enabled

Command Mode EXEC

Usage Information The output from this command is the */var/log/syslog* file.

Example

```
OS10# show trace
May 23 17:10:03 OS10 base_nas: [NETLINK:NH-
EVENT]:ds_api_linux_neigh.c:nl_to_nei
gh_info:109, Operation:Add-NH family:IPv4(2) flags:0x0 state:Failed(32) if-idx:
4
May 23 17:10:03 OS10 base_nas: [NETLINK:NH-
EVENT]:ds_api_linux_neigh.c:nl_to_nei
gh_info:120, NextHop IP:192.168.10.1
May 23 17:10:04 OS10 base_nas: [NDI:NDI-SAI], Values are invalid - can't be
conv
erted to SAI types (func:2359304)
May 23 17:10:04 OS10 base_nas: [NDI:NDI-SAI], Hash value - 20 can't be
converted
May 23 17:10:04 OS10 base_nas: [NDI:NDI-SAI], Values are invalid - can't be
conv
erted to SAI types (func:2359305)
May 23 17:10:04 OS10 base_nas: [NDI:NDI-SAI], Values are invalid - can't be
conv
erted to SAI types (func:2359311)
May 23 17:10:04 OS10 base_nas: [NDI:NDI-SAI], Hash value - 20 can't be
converted
May 23 17:10:04 OS10 base_nas: [NDI:NDI-SAI], Values are invalid - can't be
conv
erted to SAI types (func:2359312)
May 23 17:10:04 OS10 base_nas: [NDI:NDI-SAI], Invalid operation type for NDI
(23
59344)
May 23 17:10:04 OS10 base_nas: [NDI:NDI-SAI], Invalid operation type for NDI
(23
59345)
May 23 17:10:04 OS10 base_nas: [NDI:NDI-SAI], Invalid operation type for NDI
(23
59346)
May 23 17:10:04 OS10 base_nas: [NDI:NDI-SAI], Invalid operation type for NDI
(23
59319)
May 23 17:10:08 OS10 base_nas: [NETLINK:NH-
EVENT]:ds_api_linux_neigh.c:nl_to_nei
--More--
```

Supported Releases 10.2.0E or later

Log into OS10 device

Linux shell access is available for troubleshooting and diagnostic purposes only. Use `linuxadmin` for both the default user name and password. For security reasons, you must use the `system-user` command to change the default `linuxadmin` password from the command-line interface.

If you log in to the Linux shell before you change the password from the CLI, you are prompted to change password in the Linux shell. If you change the password in the Linux shell, configure the same password from the CLI to avoid inconsistent behavior. To save the new password for future logins, enter the `write memory` command.

CAUTION: Changing the system state from the Linux shell can result in undesired and unpredictable system behavior. Only use Linux shell commands to display system state and variables, or as instructed by Dell EMC Support.

```
OS10 login: linuxadmin
Password: linuxadmin      >> only for first-time login
You are required to change your password immediately (root enforced)
Changing password for linuxadmin.
(current) UNIX password: linuxadmin
Enter new UNIX password: enter a new password
Retype new UNIX password: re-enter the new password
Linux OS10 3.16.7-ckt20 #1 SMP Debian 3.16.7-ckt20-1+deb8u4 (2017-05-01) x86_64

The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.


-*-*-*
```

- Layer 2 contains information about how to configure routing information including 802.1X, LACP, LLDP, MAC, and so on.
- Layer 3 contains information about how to troubleshoot BCP, ECMP, OSPF, and so on.
- System management contains information about how to view current interface configuration information, how to view a list of all system devices, how to view the software version, and so on.
- Quality of service contains information about quality of service including classification and marking, congestion management, policing and shaping, and so on.
- Monitoring contains information about how to view alarms, events, logs, and so on.

Installation

How do I configure a default management route?

Although the default management route was configured during installation, you can use the `route add default gw` command from the Linux shell to configure the default management IP address for routing. SupportAssist requires the default management route is configured to work properly, as well as DNS configured and a route to a proxy server. For more information, see Configure SupportAssist and proxy-server.

How do I log into the OS10 shell as the system administration?

Use `linuxadmin` as the username and password to enter OS10 at root level.

Where can I find additional installation information for my specific device?

See the *Setup Guide* shipped with your device or the platform-specific *Installation Guide* on the Dell EMC Support page at dell.com/support.

Hardware

What are the default console settings for ON-Series devices?

- Set the data rate to 115200 baud
- Set the data format to 8 bits, stop bits to 1, and no parity
- Set flow control to none

How do I view the hardware inventory?

Use the `show inventory` command to view complete system inventory.

How do I view the process-related information?

Use the `show processes node-id` *`node-id-number`* `[pid` *`process-id`*`]` command to view the process CPU utilization information.

Configuration

How do I enter CONFIGURATION mode?

Use the `configure terminal` command to change from EXEC mode to CONFIGURATION mode.

I made changes to the running configuration file but the updates are not showing. How do I view my changes?

Use the `show running-configuration` command to view changes that you have made to the running-configuration file. Here are the differences between the available configuration files:

- startup-configuration contains the configuration applied at device startup
- running-configuration contains the current configuration of the device
- candidate-configuration is an intermediate temporary buffer that stores configuration changes prior to applying them to the running-configuration

Security

How do I add new users?

Use the `username` commands to add new users. Use the `show users` command to view a list of current users.

How do I view RADIUS transactions to troubleshoot problems?

Use the `debug radius` command.

How do I view the current DHCP binding information?

Use the `show ip dhcp binding` command.

Layer 2

How do I view the VLAN running configuration?

Use the `show vlan` command to view all configured VLANs.

Layer 3

How do I view IPv6 interface information?

Use the `show ipv6 route summary` command.

How do I view summary information for all IP routes?

Use the `show running-configuration` command.

How do I view summary information for the OSPF database?

Use the `show ip ospf database` command.

How do I view configuration of OSPF neighbors connected to the local router?

Use the `show ip ospf neighbor` command.

System management

How can I view the current interface configuration?

Use the `show running-configuration` command to view all currently configured interfaces.

How can I view a list of all system devices?

Use the `show inventory` command to view a complete list.

How can I view the software version?

Use the `show version` command to view the currently running software version.

Access control lists

How do I setup filters to deny or permit packets from an IPv4 or IPv6 address?

Use the `deny` or `permit` commands to create ACL filters.

How do I clear access-list counters?

Use the `clear ip access-list counters`, `clear ipv6 access-list counters`, or `clear mac access-list counters` commands.

How do I setup filters to automatically assign sequencer numbers for specific addresses?

Use the `seq deny` or `seq permit` commands for specific packet filtering.

How do I view access-list and access-group information?

Use the `show {ip | mac | ipv6} access-group` and `show {ip | mac | ipv6} access-list` commands.

Quality of service

What are the QoS error messages?

Flow control error messages:

- `Error: priority-flow-control mode is on, disable pfc mode to enable LLFC`
- `% Warning: Make sure all qos-groups are matched in a single class in attached policy-map`

Priority flow control mode error message:

```
% Error: LLFC flowcontrol is on, disable LLFC to enable PFC
```

PFC shared-buffer size error message:

```
% Error: Hardware update failed.
```

Pause error message:

```
% Error: Buffer-size should be greater than Pause threshold and Pause threshold should be
greater than equal to Resume threshold.
```

PFC cost of service error messages:

- `% Error: Not enough buffers are available, to enable system-qos wide pause for all pfc-cos values in the policymap`
- `% Error: Not enough buffers are available, to enable system-qos wide pause for the pfc-cos values in the policymap`

- `% Error: Not enough buffers are available, to enable pause for all pfc-cos values in the policymap for this interface`
- `% Warning: Not enough buffers are available, for lossy traffic. Expect lossy traffic drops, else reconfigure the pause buffers`

Monitoring

How can I check if SupportAssist is enabled?

Use the `show support-assist status` command to view current configuration information.

How can I view a list of alarms?

Use the `show alarms details` to view a list of all system alarms.

How do I enable or disable system logging?

Use the `logging enable` command or the `logging disable` command.

How do I view system logging messages?

Use the `show logging` command to view messages by log file or process name.

20

Support resources

The Dell EMC Support site provides a range of documents and tools to assist you with effectively using Dell EMC devices. Through the support site you can obtain technical information regarding Dell EMC products, access software upgrades and patches, download available management software, and manage your open cases. The Dell EMC support site provides integrated, secure access to these services.

To access the Dell EMC Support site, go to www.dell.com/support/. To display information in your language, scroll down to the bottom of the page and select your country from the drop-down menu.

- To obtain product-specific information, enter the 7-character service tag or 11-digit express service code of your switch and click **Submit**.

  To view the service tag or express service code, pull out the luggage tag on the chassis or enter the `show chassis` command from the CLI.
- To receive additional kinds of technical support, click **Contact Us**, then click **Technical Support**.

To access system documentation, see www.dell.com/manuals/.

To search for drivers and downloads, see www.dell.com/drivers/.

To participate in Dell EMC community blogs and forums, see www.dell.com/community.