NETGEAR ENTERPRISE

User Manual

Fully Managed Switches M4350 Series

202-12953-02
September 2026

Support

For product support and to access the latest product firmware, visit netgear.com/support.
For the duration of software and cyber security updates, visit netgear.com/about/eos/.

Regulatory and Legal

Si ce produit est vendu au Canada, vous pouvez accéder à ce document en français canadien à https://www.netgear.com/support/enterprise.
(If this product is sold in Canada, you can access this document in Canadian French at https://www.netgear.com/support/enterprise.)
For safety and regulatory compliance information including the EU Declaration of Conformity, visit https://www.netgear.com/about/regulatory/.
See the regulatory compliance document before connecting the power supply.
For NETGEAR's Privacy Policy, visit https://www.netgear.com/about/privacy-policy.
Where permitted by law, by using this device, you are agreeing to NETGEAR's Terms and Conditions at https://www.netgear.com/about/terms-and-conditions and if you do not agree, return the device to your place of purchase within your return period.
This product is designed and warranted for indoor use only. Do not use this device outdoors. The PoE source is intended for intra-building connection only.

Trademarks

© 2026 NETGEAR, Inc., NETGEAR, and the NETGEAR Logo are trademarks of NETGEAR, Inc.
Any non-NETGEAR trademarks are used for reference purposes only.

Contents

Get Started with the Main UI 1
Available publications and online help 1
Register your product 2
Main UI overview and other management options 2
Log in to the main UI with a web browser 3
Log in to the main UI using the switch default IP address 4
Log in to the main UI with a known IP address 5
Main UI buttons and user-defined fields 6
Interface naming conventions 7
Save your settings to the running configuration 7
Main UI online help 8
Use the Device View in the Main UI 8
Set up SNMPv3 access 10
Discover and onboard the switch to the NETGEAR Engage Controller 11
How Insight and the main UI interact with each other 12
Credentials for the main UI 13

Manage Stacking 15
Switch stacking overview 15
Stack firmware synchronization and upgrade 16
Stack configuration maintenance 16
Stack management switch election 17
Stack factory default reset behavior 17
Stacking limitations 17
Configure a stack 18
Set the stack sample mode 18
Add a switch as the stack management switch 19
Add a stack member or standby management switch 21
Change the setting for a switch in the stack 22
Remove a switch from the stack 23
Configure the stack mode for a port 24
Configure stack firmware synchronization 25
Configure the trunk hash mode 26
Display information about the stack configuration 27

Copyright © 2026 NETGEAR

Display the heartbeat message stack status 28
Display stack port diagnostic, status, and path information 29
Remove a stack 30
Stack nonstop forwarding 30
Configure nonstop forwarding and display associated information 31
Display the nonstop forwarding checkpoint statistics 34

Configure Switch System Information 35
Switch system and slot information 35
View and configure switch system information 35
View the fan status 38
View the temperature sensor information 39
View the device status and firmware version 40
View the system CPU status 41
Configure the CPU thresholds 42
View or clear switch statistics 43
View USB device information 45
Display information about a stack and optional switch models for a stack 46
Loopback interface 48
IPv4 management interfaces and VLANs 49
Configure the IPv4 service port 49
Configure an IPv4 management VLAN 51
Configure an IPv4 management interface 53
IPv6 management interfaces and VLANs 54
Configure an IPv6 service port 55
Manage IPv6 addresses for the IPv6 service port 56
Configure an IPv6 management VLAN 57
Manage IPv6 addresses for the IPv6 management VLAN 60
Manage the IPv6 default route addresses for the IPv6 management VLAN 60
Configure an IPv6 management interface 61
Manage IPv6 addresses for the IPv6 management interface 64
Manage the IPv6 default route addresses for the IPv6 management interface 64
Time and NTP settings 65
Configure the global NTP settings 65
Display the global NTP status and packet information 67
Configure NTP servers 69
Configure daylight saving time settings 72
View the daylight saving time status 74
Precision Time Protocol transparent clock 75
Manage the global PTP TC settings 76

Copyright © 2026 NETGEAR

Manage the PTP TC interface settings . . . 76
Manage PTP multiple domains . . . 77
Precision Time Protocol boundary clock and grandmaster clock . . . 78
Manage the PTP global configuration settings . . . 79
Manage the PTP domain, boundary clock, and grandmaster clock settings . . . 80
Manage the PTP domain, boundary clock, and grandmaster interface settings . . . 83
View or clear the PTP boundary clock statistics . . . 84
View detailed information about the PTP boundary clock . . . 86
Domain Name System . . . 88
Configure the global DNS settings and add a DNS server . . . 89
Remove a DNS server . . . 90
Configure and view host name-to-IP address information . . . 91
Switch database management template . . . 93
Green Ethernet settings . . . 94
Configure the global green Ethernet settings . . . 94
Configure green Ethernet interface settings . . . 95
Configure and display detailed green Ethernet settings for an interface . . . 96
Display green Ethernet information for the link partner of an interface . . . 99
Display the green Ethernet statistics summary . . . 99
Bonjour settings . . . 102
Enable or disable Bonjour . . . 102
Display Bonjour information . . . 102
Enable or disable the TFTP server capability . . . 103
IPv4 DHCP server . . . 103
Configure a DHCP server . . . 103
Add a DHCP class for use with DHCP Option 82 . . . 105
Configure a DHCP class . . . 105
Manage DHCP pools . . . 106
Configure DHCP pool options . . . 109
Bind a DHCP pool to a class and configure the binding . . . 110
Display DHCP server statistics . . . 111
Display the DHCP bindings . . . 112
Delete one or all dynamic DHCP bindings . . . 113
View bindings with DHCP conflicts . . . 114
Delete one or all DHCP bindings with conflicts . . . 114
DHCP relay . . . 115
Configure the global DHCP relay settings and display the relay statistics . . . 115
Configure a DHCP relay interface . . . 118
DHCP Layer 2 relay . . . 119
Configure the global DHCP L2 relay settings . . . 119

Copyright © 2026 NETGEAR

Configure a DHCP L2 relay interface . . . 120
Display DHCP L2 relay interface statistics . . . 121
UDP relay . . . 122
Configure the global UDP relay settings and add a UDP relay . . . 122
Change a UDP relay configuration . . . 123
Remove a UDP relay configuration . . . 124
Add a UDP interface configuration . . . 124
Change a UDP interface configuration . . . 125
Remove a UDP interface . . . 126
DHCPv6 server . . . 126
Enable the DHCPv6 server . . . 126
Manage DHCPv6 pools . . . 127
Manage DHCPv6 prefix delegation for pools . . . 129
Configure the DHCPv6 settings for an interface . . . 131
Display the DHCPv6 bindings . . . 132
Display DHCPv6 server statistics . . . 133
Delete DHCPv6 statistics for one or all interfaces . . . 135
DHCPv6 relay interface . . . 136
Power over Ethernet . . . 137
PoE concepts . . . 137
Set the PoE system usage threshold and power management mode . . . 141
Configure the PoE ports settings . . . 142
Power-cycle one or more PoE ports . . . 148
Manage the PoE usage threshold . . . 148
Manage N+1 power redundancy . . . 150
Display information about multiple power source management for PoE power . . . 151
Timer schedules . . . 153
Create a timer schedule . . . 153
Specify the settings for an absolute timer schedule . . . 154
Specify the settings for a recurring timer schedule . . . 155
Change the settings for a recurring timer schedule entry . . . 157
Delete a timer schedule entry . . . 157
Delete a timer schedule . . . 158
Simple Network Management Protocol . . . 158
Manage SNMPv1 and SNMPv2 communities . . . 158
Manage the SNMPv1 and SNMPv2 trap settings . . . 160
Configure SNMPv1 and SNMPv2 trap flags . . . 163
Display the supported MIBs . . . 168
Manage SNMPv3 users . . . 168
Link Layer Discovery Protocol . . . 170

Copyright © 2026 NETGEAR

Configure the global LLDP settings . . . 170
Configure LLDP interface settings . . . 171
Display or clear LLDP statistics . . . 172
Display LLDP local device information . . . 174
Display LLDP remote device information . . . 175
Display the LLDP remote device inventory . . . 176
Link Layer Discovery Protocol for Media Endpoint Devices . . . 177
Configure the global LLDP-MED settings . . . 177
Configure LLDP-MED interface settings . . . 178
Display LLDP-MED local device information . . . 179
Display the LLDP-MED remote device information . . . 182
Display the LLDP-MED remote device inventory . . . 185
Link dependency . . . 185
Configure a link dependency group . . . 186
Configure or display upstream and downstream interfaces for a link dependency group . . . 187
Clear all interfaces in a link dependency group . . . 188
Industry Standard Discovery Protocol . . . 189
Configure the global ISDP settings . . . 189
Configure ISDP settings for an interface . . . 190
Display or clear ISDP neighbor information . . . 191
Display or clear ISDP statistics . . . 192
Configure the RGB LED color settings . . . 193

Configure Switching Information . . . 195
VLANs . . . 195
Manage the VLAN configuration on the switch . . . 195
Auto-Trunk overview . . . 199
Enable or disable Auto-Trunks . . . 200
Configure the switch port mode settings for interfaces . . . 201
Configure membership interfaces for a VLAN . . . 203
View the VLAN status on the switch . . . 204
Change the port VLAN ID settings . . . 205
Configure a MAC-based VLAN . . . 207
Configure an OUI-based VLAN . . . 208
Configure a protocol-based VLAN group . . . 210
Configure membership interfaces for a protocol-based VLAN group . . . 212
Configure an IP subnet-based VLAN . . . 213
Configure a double VLAN . . . 214
Configure a voice VLAN . . . 216

Copyright © 2026 NETGEAR

Configure Generic Attribute Registration Protocol . . . 217
Auto-VoIP . . . 220
Configure Auto-VoIP protocol-based settings . . . 220
Configure the Auto-VoIP OUI-based properties . . . 222
Configure the OUI-based interface settings . . . 222
Manage the OUI table . . . 223
Display the Auto-VoIP status . . . 225
Internet Small Computer System Interface . . . 225
Enable iSCSI and configure the QoS settings for iSCSI traffic . . . 226
Add an iSCSI target . . . 227
View iSCSI sessions . . . 227
View detailed information about iSCSI sessions . . . 228
Spanning Tree Protocol . . . 229
Configure the STP settings and display the STP status . . . 230
Configure the CST settings and display the CST status . . . 232
Configure the CST interface settings . . . 235
Display the CST interface status . . . 237
Manage MST instances . . . 240
Configure and display the interface settings for an MST instance . . . 242
Display the STP interface statistics . . . 245
Configure the PVST/RPVST VLAN settings . . . 246
Change a PVST/RPVST VLAN configuration . . . 247
Remove a PVST/RPVST VLAN configuration . . . 248
Configure the PVST and RPVST interface settings . . . 248
Display the PVST statistics . . . 250
Multicast forwarding database . . . 251
Display the entries in the multicast forwarding database . . . 251
Remove the IGMP snooping entries from the multicast forwarding database . . . 252
Remove all known multicast MAC entries from the multicast forwarding database . . . 253
Display the multicast forwarding database statistics . . . 253
Internet Group Management Protocol snooping . . . 254
Configure IGMP snooping automatically with IGMP Plus mode . . . 255
Configure IGMP snooping manually . . . 256
Configure the IGMP snooping settings for interfaces . . . 258
Configure IGMP snooping for VLANs automatically with IGMP Plus mode . . . 260
Configure IGMP snooping for VLANs manually . . . 261
Configure an IGMP multicast router interface . . . 263
Configure an IGMP multicast router VLAN . . . 264
IGMP snooping querier overview . . . 264
Configure the IGMP snooping querier global settings . . . 264

Copyright © 2026 NETGEAR

Configure an IGMP snooping querier for a VLAN . . . 265
Remove the IGMP snooping querier settings for a VLAN . . . 267
Display the status of the IGMP snooping querier . . . 268
Multicast Listener Discovery snooping . . . 268
Configure MLD snooping automatically with MLD Plus mode . . . 269
Configure MLD snooping manually . . . 270
Configure the MLD snooping settings for interfaces . . . 271
Configure MLD snooping for VLANs automatically with MLD Plus mode . . . 273
Configure MLD snooping for VLANs manually . . . 274
Remove the MLD snooping querier settings for a VLAN . . . 275
Configure an MLD multicast router interface . . . 276
Configure an MLD multicast router VLAN . . . 277
Configure the MLD snooping querier global settings . . . 277
Configure an MLD snooping querier for a VLAN . . . 278
Remove the MLD snooping querier settings for a VLAN . . . 280
Multicast VLAN registration . . . 281
Configure the global MVR settings . . . 281
Configure an MVR group . . . 282
Remove an MVR group . . . 283
Configure an MVR interface . . . 283
Configure the interface members of an MVR group . . . 285
Display the MVR statistics . . . 285
MAC address table . . . 286
View, search, or clear the MAC address table . . . 286
Set the dynamic address aging interval . . . 287
Add a static MAC address to the MAC address table . . . 288
Remove a static MAC address from the MAC address table . . . 288
Port settings . . . 289
Configure and display the port settings . . . 289
Add port, LAG, and VLAN descriptions . . . 293
Display transceiver module information . . . 295
Configure the expandable port settings . . . 296
Configure the port link flap settings . . . 297
Link aggregation groups . . . 298
Auto-LAG overview . . . 298
Enable or disable Auto-LAGs . . . 299
Enable or disable the Auto-LAG global AVB profile mode . . . 299
Configure the hash mode for Auto-LAGs . . . 300
Configure the non-unicast global hash mode . . . 301
Configure the settings for a LAG . . . 302

Copyright © 2026 NETGEAR

Configure a single LAG and its membership . . . 305
MLAG overview . . . 306
Configure the basic VPC global settings . . . 308
Configure the advanced VPC global settings . . . 313
802.1AS timing and synchronization . . . 323
Configure and view the global 802.1AS settings . . . 323
Configure the 802.1AS interface settings . . . 324
View the 802.1AS statistics . . . 327
Multiple Registration Protocol and 802.1Qav . . . 329
Configure the global MRP settings . . . 330
Configure 802.1Qav mapping . . . 331
Configure the MRP interface settings . . . 332
Display or clear MMRP statistics . . . 334
Display or clear MVRP statistics . . . 335
Display or clear MSRP statistics . . . 337
Display the MSRP reservation settings . . . 339
Configure and display the Qav settings for interfaces . . . 342
Display MSRP streams information . . . 343
Loop protection . . . 345
About loop protection . . . 346
Configure the global loop protection settings . . . 347
Configure the loop protection settings for interfaces and display the loop protection state . . . 348

Manage Routing . . . 350
Routing concepts . . . 350
Routing table, routes and route preferences . . . 350
Configure a route and display learned routes . . . 350
Delete a route . . . 352
Specify route preferences . . . 353
IPv4 routing . . . 353
Manage the global IPv4 routing settings . . . 354
Display the IPv4 statistics . . . 355
Configure IPv4 routing interfaces . . . 359
Delete the routing IP address from an IPv4 routing interface . . . 361
Configure a secondary IP address for an IPv4 routing interface . . . 362
Delete the secondary IP address from an IPv4 routing interface . . . 363
IPv6 routing . . . 363
Manage the global IPv6 routing settings . . . 363
Display the IPv6 route table . . . 364

Copyright © 2026 NETGEAR

Configure IPv6 routing interfaces . . . 365
Configure prefix settings for an IPv6 routing interface . . . 368
Delete a prefix setting from an IPv6 routing interface . . . 369
Display the IPv6 and ICMPv6 statistics for an IPv6 routing interface . . . 370
Display the IPv6 neighbor table or clear IPv6 neighbor entries . . . 375
Configure IPv6 static routes . . . 377
Delete an IPv6 static route . . . 378
Configure the IPv6 route preference for the switch . . . 378
Routing VLANs . . . 379
Create a routing VLAN with the VLAN static routing wizard . . . 379
Configure routing for an existing VLAN . . . 381
Remove the routing function from a VLAN . . . 381
Address Resolution Protocol . . . 382
Display the ARP entries in the ARP cache . . . 383
Add or change a static entry in the ARP table . . . 383
Delete a static ARP entry . . . 385
Configure the ARP table settings or remove entries from the table . . . 385
Routing Information Protocol . . . 386
Enable or disable RIP on the switch . . . 387
Configure the global RIP settings for the switch . . . 387
Configure RIP interface settings . . . 389
Configure the RIP route redistribution settings and display the route redistribution summary . . . 391
Router discovery and router advertisements . . . 393
Virtual Router Redundancy Protocol . . . 395
Enable VRRP and add a primary virtual router . . . 395
Enable VRRP and add a primary virtual router with enhanced settings . . . 396
Change the settings for a primary virtual router . . . 399
Remove a primary virtual router . . . 400
Configure a secondary virtual router . . . 400
Change the IP address for a secondary virtual router . . . 401
Remove a secondary virtual router . . . 401
Configure VRRP interface tracking and route tracking . . . 402
Display VRRP statistics . . . 403

Configure OSPF and OSPFv3 . . . 406
Open Shortest Path First . . . 406
Enable OSPF . . . 407
Configure the OSPF default route advertisement . . . 407
Configure the global OSPF settings . . . 408

Copyright © 2026 NETGEAR

Add or delete an OSPF common area ID . . . 411
Add an OSPF stub area . . . 412
Add an OSPF NSSA area . . . 414
Add an OSPF area range . . . 416
Configure an OSPF interface . . . 417
View or clear OSPF statistics for an interface . . . 421
View or clear OSPF neighbor information for an interface . . . 424
View the OSPF link state database . . . 426
Configure an OSPF virtual link . . . 430
Configure the OSPF route redistribution . . . 433
Configure OSPF nonstop forwarding . . . 435
Open Shortest Path First version 3 . . . 436
Enable OSPFv3 . . . 436
Configure the OSPFv3 default route advertisement . . . 437
Configure the global OSPFv3 settings . . . 438
Add or delete an OSPFv3 common area ID . . . 440
Add an OSPFv3 stub area . . . 442
Add an OSPFv3 NSSA area . . . 443
Add an OSPFv3 area range . . . 445
Configure an OSPFv3 interface . . . 446
View or clear OSPFv3 statistics for an interface . . . 450
View or clear OSPFv3 neighbor information for an interface . . . 452
View the OSPFv3 link state database . . . 455
Configure an OSPFv3 virtual link . . . 459
Configure the OSPFv3 route redistribution . . . 462
Configure OSPFv3 nonstop forwarding . . . 463

Configure Multicast Routing . . . 465
IPv4 multicast routing and the IPv4 multicast route table . . . 465
Display the IPv4 multicast route table . . . 465
Add static multicast entries to the IPv4 Mroute table . . . 466
Delete a static multicast entry from the IPv4 Mroute table . . . 467
Configure global multicast settings for the switch . . . 467
Configure a multicast interface . . . 468
Distance Vector Multicast Routing Protocol . . . 469
Enable DVMRP on the switch and view route information . . . 470
Configure a DVMRP interface . . . 470
View DVMRP neighbors . . . 472
View the DVMRP next hops . . . 473
View the DVMRP prune table . . . 474

Copyright © 2026 NETGEAR

View the DVMRP routes . . . 475
IGMP for IPv4 multicast routing . . . 476
Enable or disable IGMP for the switch . . . 476
Configure an IGMP routing interface . . . 477
Display the statistics for the IGMP routing interfaces . . . 479
Display the IGMP groups and search the IGMP group database . . . 480
Display the IGMP membership information and search the IGMP membership database . . . 482
Configure an IGMP proxy interface . . . 484
Display the statistics for the IGMP proxy interface . . . 485
Display the IGMP proxy membership and search the IGMP proxy membership database . . . 486
PIM for IPv4 multicast routing . . . 487
Configure the global PIM IPv4 settings on the switch . . . 488
Add IPv4 PIM-SSM groups . . . 488
Delete an IPv4 PIM-SSM group . . . 489
Configure an IPv4 PIM interface . . . 489
Display IPv4 PIM neighbors and search the PIM neighbor database . . . 491
Add an IPv4 PIM candidate rendezvous point configuration . . . 492
Delete an IPv4 PIM candidate rendezvous point configuration . . . 493
Configure an interface as an IPv4 PIM bootstrap router candidate . . . 493
Delete an IPv4 PIM bootstrap router candidate configuration . . . 494
Configure a static IPv4 PIM rendezvous point for a group . . . 494
Delete a static IPv4 PIM rendezvous point configuration . . . 495
Static multicast routes for IPv4 addresses . . . 496
Configure static multicast routes for IPv4 addresses . . . 496
Delete a static multicast route for an IPv4 address . . . 496
Multicast admin boundaries for IPv4 addresses . . . 497
Configure an interface as a multicast admin boundary . . . 497
Delete a multicast admin boundary configuration for an interface . . . 498
IPv6 multicast routing and the IPv6 multicast route table . . . 498
Display the IPv6 multicast route table . . . 498
PIM for IPv6 multicast routing . . . 499
Configure the global PIM IPv6 settings on the switch . . . 500
Add IPv6 PIM-SSM groups . . . 500
Delete an IPv6 PIM-SSM group . . . 501
Configure an IPv6 PIM interface . . . 501
Display IPv6 PIM neighbors and search the PIM neighbor database . . . 503
Add an IPv6 PIM candidate rendezvous point configuration . . . 504
Delete an IPv6 PIM candidate rendezvous point configuration . . . 505

Copyright © 2026 NETGEAR

Configure an interface as an IPv6 PIM bootstrap router candidate . . . 505
Delete an IPv6 PIM bootstrap router candidate configuration . . . 506
Configure a static IPv6 PIM rendezvous point for a group . . . 506
Delete a static IPv6 PIM rendezvous point configuration . . . 507
MLD for IPv6 multicast routing . . . 508
Configure the global MLP settings for the switch . . . 508
Configure an MLD routing interface . . . 508
Display the statistics for the MLD routing interfaces . . . 510
Display the MLD groups and search the MLD group database . . . 511
Display or clear MLD traffic statistics . . . 513
Configure an MLD proxy interface . . . 513
Display the statistics for the MLD proxy interface . . . 515
Display the MLD proxy membership and search the MLD proxy membership database . . . 516
Static multicast routes for IPv6 addresses . . . 517
Configure static multicast routes for IPv6 addresses . . . 517
Delete a static multicast route for an IPv6 address . . . 518

Configure Quality of Service . . . 520
Quality of Service concepts . . . 520
Class of Service . . . 520
CoS configuration concepts . . . 520
Configure the CoS trust mode settings globally or for a specific interface . . . 521
Map 802.1p priorities to queues . . . 522
Map DSCP values to queues . . . 523
Configure the CoS interface settings for an interface . . . 524
Configure CoS queue settings for an interface . . . 525
Configure the CoS WRED precedence settings for dropping packets . . . 527
Differentiated Services . . . 529
Defining DiffServ . . . 529
DiffServ wizard overview . . . 530
Use the DiffServ wizard to create a traffic class and policy for one or more interfaces . . . 530
Configure the DiffServ mode and display the entries in the DiffServ private MIB tables . . . 531
Configure a DiffServ class . . . 532
Configure an IPv6 DiffServ class . . . 538
Configure a DiffServ policy . . . 543
Configure the DiffServ service interface . . . 550
Display DiffServ service statistics . . . 553

Copyright © 2026 NETGEAR

Manage Switch Security 555
User accounts and passwords 555
Add or change a user account 556
Delete a user account 557
Configure user password requirements 558
Enable multi-factor authentication on the switch 559
Change the privileged EXEC CLI password 560
Change the console, Telnet, or SSH password 560
RADIUS servers 562
Configure the global RADIUS server settings 562
Configure a RADIUS authentication server on the switch 564
Configure a RADIUS accounting server on the switch 567
TACACS+ servers 570
Configure the global TACACS+ settings 571
Add a TACACS+ server to the switch 572
Modify the settings for a TACACS+ server on the switch 572
Remove a TACACS+ server from the switch 573
Authentication lists 573
Configure a login authentication list 574
Delete a login authentication list 575
Configure an enable authentication list 576
Delete an enable authentication list 577
Configure the Dot1x authentication list 578
Configure the HTTP authentication list 578
Configure the HTTPS authentication List 580
Current login sessions 581
HTTP and HTTPS management access 581
Configure the HTTP access settings 581
Configure the HTTPS access settings 582
Browser security message with HTTPS access 583
Manage certificates for HTTPS access 584
SSH management access 588
Configure the global SSH access settings 588
Manage RSA, DSA, and ECDSA keys for SSH access 590
Telnet management access 593
Select Telnet authentication lists 594
Configure inbound Telnet settings 594
Configure outbound Telnet settings 595
Console port management access 596

Copyright © 2026 NETGEAR

Denial of service . . . 597
Management access profiles and rules . . . 599
Add an access profile . . . 600
Add a rule to the access profile . . . 600
Activate the access profile . . . 601
Display the access profile summary and the number of filtered packets . . . 602
Deactivate an access profile . . . 602
Remove an access profile . . . 603
Port authentication . . . 603
Configure the global 802.1X authentication settings . . . 604
Manage port authentication on individual ports . . . 606
Display the port summary . . . 612
Display the client summary . . . 614
MAC filters for traffic control . . . 615
Create a MAC filter . . . 616
Delete a MAC filter . . . 617
Display the MAC filter summary . . . 618
Port security . . . 618
Configure the global port security mode . . . 619
Configure a port security interface . . . 620
Display learned MAC addresses and convert them to static addresses . . . 621
Add a static MAC address to the MAC address table for port security . . . 622
Remove a static MAC address from the MAC address table for port security . . . 623
Private port groups . . . 623
Add a private port group . . . 623
Remove a private port group . . . 624
Configure the membership of a private port group . . . 624
Protect ports . . . 625
Private VLANs . . . 626
Overview of the tasks for private VLAN configuration . . . 627
Assign a private VLAN type to a VLAN . . . 628
Configure a private VLAN association with a primary and secondary VLAN . . . 629
Remove an existing private VLAN association . . . 630
Configure the private VLAN port mode . . . 631
Private VLAN host interface: Assign the interface to primary and secondary VLANs . . . 633
Private VLAN host interface: Remove the interface from primary and secondary VLANs . . . 634
Private VLAN promiscuous interface: Assign the interface to primary and secondary VLANs . . . 635
Private VLAN promiscuous interface: Remove the interface from primary and secondary

Copyright © 2026 NETGEAR

VLANs . . . 636
Private VLAN promiscuous trunk interface: Add primary and secondary VLANs to the trunk . . . 637
Private VLAN promiscuous trunk interface: Remove primary and secondary VLANs from the trunk . . . 638
Private VLAN isolated trunk interface: Add primary and secondary VLANs to the trunk . . . 639
Private VLAN isolated trunk interface: Remove primary and secondary VLANs from the trunk . . . 640
Configure native and allowed VLANs on a private VLAN trunk interface . . . 641
Storm control . . . 642
Configure global storm control settings . . . 643
Configure storm control settings for one or more ports . . . 644
DHCP snooping . . . 646
Enable DHCP snooping for the switch . . . 646
Enable DHCP snooping for a VLAN . . . 647
Configure DHCP snooping interface settings . . . 648
Add a static DHCP binding and display or clear dynamic DHCP bindings . . . 649
Remove a static DHCP binding . . . 650
Configure DHCP snooping persistent settings . . . 651
Display or clear DHCP snooping statistics . . . 651
DHCPv6 snooping . . . 652
Enable DHCPv6 snooping for the switch . . . 653
Enable DHCPv6 snooping for a VLAN . . . 653
Configure DHCPv6 snooping interface settings . . . 654
Add a static DHCPv6 binding and display or clear dynamic DHCPv6 bindings . . . 656
Remove a static DHCPv6 binding . . . 657
Configure DHCPv6 snooping persistent settings . . . 657
Display or clear DHCPv6 snooping statistics . . . 658
IP source guard interfaces . . . 659
Configure IP source guard on an interface . . . 659
Add a static IP source guard binding and display or clear dynamic IP source guard bindings . . . 661
Remove a static IP source guard binding . . . 662
IPv6 source guard interfaces . . . 662
Configure IPv6 source guard on an interface . . . 663
Add a static IPv6 source guard binding and display or clear dynamic IPv6 source guard bindings . . . 664
Remove a static IPv6 source guard binding . . . 665
Dynamic ARP inspection . . . 666

Copyright © 2026 NETGEAR

Configure the global DAI settings . . . 666
Configure DAI VLANs . . . 667
Configure DAI interfaces . . . 668
Create a DAI access control list . . . 669
Configure a rule for an existing DAI ACL . . . 669
Delete a rule from an existing DAI ACL . . . 670
Delete a DAI access control list . . . 670
Display the DAI statistics . . . 671
Captive portals . . . 672
Configure the global captive portal settings . . . 673
Configure a captive portal . . . 675
Delete a captive portal . . . 677
Configure a captive portal binding . . . 678
Display or delete captive portal bindings in the captive portal binding table . . . 678
Configure captive portal groups . . . 679
Configure captive portal users . . . 680
Configure the captive portal trap flag settings . . . 683
Display or clear captive portal client statistics . . . 684
Access control lists . . . 685
Use the ACL Wizard to create a simple ACL . . . 686
Configure a MAC ACL . . . 691
Configure MAC ACL rules . . . 693
Configure MAC bindings . . . 697
Display or delete MAC ACL bindings in the MAC binding table . . . 698
Configure a basic or extended IPv4 ACL . . . 699
Configure rules for a basic IP ACL . . . 701
Configure rules for an extended IPv4 ACL . . . 705
Configure an IPv6 ACL . . . 713
Configure rules for an IPv6 ACL . . . 715
Configure IP ACL interface bindings . . . 720
Display or delete IP ACL bindings in the IP ACL binding table . . . 722
Configure VLAN ACL bindings . . . 722

Monitor the Switch and Network . . . 725
Port and EAP packet statistics . . . 725
Display or clear port statistics . . . 725
Display or clear detailed statistics for a port . . . 727
Display or clear EAP and EAPoL statistics . . . 736
Perform a cable test . . . 738
Logs . . . 739

Copyright © 2026 NETGEAR

Manage and display the memory log . . . 739
Message log format . . . 741
Enable or disable the command configuration log . . . 742
Enable or disable console logging . . . 742
Syslog and log server host settings . . . 743
Configure the syslog settings . . . 743
Add a syslog server . . . 745
Modify the settings for a syslog server . . . 746
Delete the settings for a syslog server . . . 747
Trap log . . . 747
Event log . . . 748
Configure USB logging . . . 749
Port mirroring . . . 750
RSPAN VLANs and source and destination switches . . . 752
Configure an existing VLAN as an RSPAN VLAN . . . 753
Configure the switch as an RSPAN source switch . . . 753
Configure the switch as an RSPAN destination switch . . . 756
sFlow monitoring . . . 757
sFlow agent overview . . . 757
Configure the source interface for the sFlow agent . . . 757
Configure an sFlow receiver . . . 758
Configure sFlow polling and sampling on an interface . . . 760

Maintenance and Troubleshooting . . . 762
Save the configuration . . . 762
Automatic installation of the configuration file . . . 762
Configure the auto install process . . . 763
Option 125 DHCP server requirements for obtaining an configuration file through auto install . . . 765
Reboot the switch from the main UI . . . 766
Reset the switch to the factory default settings . . . 767
Export a file from the switch . . . 767
Export a file from the switch to a server . . . 767
Use HTTP to export a file from the switch to a computer . . . 769
Export a file from the switch to a USB storage device . . . 770
Update software or download a file . . . 771
Download a software file or another type of file from a server to the switch . . . 772
Use HTTP to download a software file or another type of file to the switch . . . 775
Download a software file or another type of file from a USB storage device to the switch . . . 777

Copyright © 2026 NETGEAR

Download and install an SSL security certificate file on the switch . . . 779
Manage software images . . . 780
Copy a software image . . . 780
Configure dual image settings . . . 781
Diagnostics and troubleshooting . . . 782
Ping an IPv4 address . . . 783
Ping an IPv6 address . . . 784
Send an IPv4 traceroute . . . 786
Send an IPv6 traceroute . . . 787
Capture Packets . . . 788
Perform a full memory dump . . . 789

Configuration Examples . . . 791
Virtual Local Area Networks (VLANs) . . . 791
VLAN configuration examples . . . 792
Access control lists (ACLs) . . . 793
MAC ACL sample configuration . . . 794
Basic IP ACL sample configuration . . . 795
Differentiated Services (DiffServ) . . . 796
Class . . . 796
DiffServ traffic classes . . . 797
Create policies . . . 797
DiffServ example configuration . . . 799
802.1X port access control . . . 800
802.1X example configuration . . . 802
Multiple Spanning Tree Protocol . . . 803
MSTP example configuration . . . 805
VLAN routing interfaces . . . 807

Software Default Settings and Hardware Specifications . . . 809
Access default settings for the switch main UI . . . 809
System features default settings . . . 810
Switching features default settings . . . 816
Routing, OSPF, OSPFv3, and multicast features default settings . . . 826
QoS features default settings . . . 834
Security features default settings . . . 835
Monitoring features default settings . . . 844
General hardware technical specifications . . . 846
Model-specific hardware technical specifications . . . 847
M4350-8X8F (SKU XSM4316) . . . 847

Copyright © 2026 NETGEAR

M4350-12X12F (SKU XSM4324) 847
M4350-24G4XF (SKU GSM4328) 848
M4350-48G4XF (SKU GSM4352) 848
M4350-24X4V (SKU XSM4328CV) 849
M4350-24F4V (SKU XSM4328FV) 849
M4350-44M4X4V (SKU MSM4352) 850
M4350-36X4V (SKU XSM4340CV) 850
M4350-24X8F8V (SKU XSM4340V) 851
M4350-32F8V (SKU XSM4340FV) 851
M4350-16V4C (SKU VSM4320C) 852
M4350-40X4C (SKU XSM4344C) 852
M4350-8M2V (SKU MSM4310) 853
M4350-24M4X4V (SKU MSM4332) 853
M4350-24F4X (SKU MSM4328F) 854
M4350-40F4C (SKU XSM4344FC) 854
M4350-16M4V (SKU MSM4320) 855
M4350-16C (SKU CSM4316) 855

Copyright © 2026 NETGEAR

Get Started with the Main UI

This user manual is for the NETGEAR Fully Managed Switches M4350 Series and covers all M4350 switch models.
This chapter provides an overview of how you can using your switch and access the main user interface (UI), also referred to as the *main UI*.
The chapter contains the following sections:

- Available publications and online help
- Register your product
- Main UI overview and other management options
- Log in to the main UI with a web browser
- Use the Device View in the Main UI
- Set up SNMPv3 access
- Discover and onboard the switch to the NETGEAR Engage Controller
- How Insight and the main UI interact with each other
- Credentials for the main UI

NOTE: For more information about the topics covered in this manual, visit the support website at netgear.com/support.

NOTE: Firmware updates with new features and bug fixes are made available from time to time at netgear.com/support/download/. Some products can regularly check the site and download new firmware, or you can check for and download new firmware manually. If the features or behavior of your product does not match what is described in this guide, you might need to update your firmware.

Available publications and online help

You can download the following publications and more for the NETGEAR Fully Managed Switches M4350 Series by visiting netgear.com/support/download/.

- Installation Guide
- Hardware Installation Guide
- Main User Manual (this manual)
- Audio Video User Manual
- Application Notes: How to stack NETGEAR M4300 switches, which also applies to the M4350 switches
- Frequently Asked Questions
- Data sheet
- For a list of supported M4350 Series models, see kb.netgear.com/000066110.

When you are logged in to the main UI, you can access documentation online by selecting **Help > Online Help > User Guide**. See Main UI online help on page 8.
In addition, context-sensitive online help is available in the main UI.

Register your product

To qualify for product updates and product warranty, we encourage you to register your product. The first time that you log in to the switch, you can register with NETGEAR by clicking the **REGISTER** button.
Registration confirms that your email alerts work, lowers technical support resolution time, and ensures the accuracy of your shipping address. We would also like to incorporate your feedback into future product development. We never sell or rent your email address and you can opt out of communications.

To register your switch with NETGEAR:

1. Visit the NETGEAR website for registration at https://my.netgear.com/registration/login.aspx.
2. Click the **Login** button, and follow the directions onscreen to register the switch with your NETGEAR email address and password.
   If you did not yet create a NETGEAR account, click the **Create account** link, follow the directions onscreen to create an account, and then register the switch with your NETGEAR email address and password.

Main UI overview and other management options

Your switch contains an embedded web server and management software for managing and monitoring switch functions. The switch functions as a simple switch without the management software. However, you can use the management software to configure more advanced features that can improve switch efficiency and overall network performance.
The switch software includes a set of comprehensive management functions for configuring and monitoring the system by using one of the following methods:

- **Main user interface (main UI)**: The main UI is a web-based management tool that lets you monitor, configure, and control your switch remotely using a standard web browser. You can access the main UI either over an Ethernet network port or over the out-of-band (OOB) port (also referred to as the service port). From your web browser, you can monitor the performance of your switch and optimize its configuration for your network. You can configure all available switch features, such as VLANs, QoS, and ACLs, by using the main UI. This manual describes how to use the main UI to manage and monitor the switch.
- **Audio-video user interface (AV UI)**: The AV UI is a web-based management tool that lets you configure and apply network profiles to sites and configure many advanced features that

can improve AV flows, switch efficiency, and overall network performance. You can access the AV UI either over an Ethernet network port or over the OOB port. For more information, see the *Audio Video User Manual*, which you can download by visiting netgear.com/support/download and entering your model number in the search box.

- **NETGEAR Engage Controller application**: The NETGEAR Engage Controller application lets you discover, onboard, manage, and monitor the switch in an AV network or other type of business network. For more information, see Discover and onboard the switch to the NETGEAR Engage Controller on page 11, Credentials for the main UI on page 13, and the *Engage Controller User Manual*, which you can download by visiting netgear.com/support/enterprise and entering your model number in the search box.
- **NETGEAR Insight**: The NETGEAR Insight Cloud Portal and Insight app let you discover, onboard, manage, and monitor the switch in an Insight network.

  NOTE: NETGEAR Insight and the Engage Controller application are mutually exclusive. To change the management mode of the switch from Insight to the Engage Controller, or the other way around, reset the switch to factory default settings.

  For information about NETGEAR Insight subscriptions, visit netgear.com/business/services/insight/subscription and kb.netgear.com/000061848/.
  This manual does not describe NETGEAR Insight procedures, which are documented in the NETGEAR knowledge base. For knowledge base articles about NETGEAR Insight, visit https://kb.netgear.com/000044336.
- **Simple Network Management Protocol (SNMP)**: SNMP is enabled by default. The switch uses both standard public MIBs for standard functionality and private MIBs that support additional switch functionality. All private MIBs begin with a "-" prefix. The main object for interface configuration is in -SWITCHING-MIB, which is a private MIB. Some interface configurations also involve objects in the public MIB, IF-MIB.
- **Command-line interface (CLI)**: The CLI is a text-based way to manage and monitor the switch. You can access the CLI by using a direct serial connection, or by using a remote logical connection with telnet or SSH. For more information about the CLI, see the *CLI Command Reference Manual*, which you can download in the near future by visiting netgear.com/support/enterprise and entering your model number in the search box.

Each of the standards-based management methods allows you to configure and monitor the components of the switch. The method you use to manage the system depends on your network size and requirements, and on your preference.

Log in to the main UI with a web browser

You can use a web browser to access the switch and log in. You must be able to ping the IP address of the management interface (any Ethernet network port) or out-of-band (OOB) port

from your computer for web access to be available:

- **Management interface**: By default, the DHCP client of the management interface is enabled so that the interface can receive an IP address from a DHCP server in your network. If the management interface does not receive an IP address from a DHCP server in your network, the IP address for the interface is set to 169.254.100.100 with 255.255.0.0 as the subnet mask. The same occurs if you connect the management interface directly to a computer and reboot the switch.
- **OOB port**: By default, the DHCP client of the OOB port is enabled so that the port can receive an IP address from a DHCP server in your network. If the OOB port does not receive an IP address from a DHCP server in your network, the IP address for the port is set to 192.168.0.239 with 255.255.0.0 as the subnet mask. The same occurs if you connect the OOB port directly to a computer and reboot the switch.

If you let a DHCP server in your network assign an IP address to switch, determine the IP address by accessing the DHCP server or by using an IP scanner utility.
For more information about logging in to the switch for the first time, see the information in the installation guide.

NOTE: The first time that you log in as an admin user to the main UI, no password is required (that is, the password is blank). After you log in for the first time, you are required to specify a local device password that you must use each subsequent time that you log in. (You can change the password again.)

Log in to the main UI using the switch default IP address

To use the switch default IP address to access the switch over the main UI:

1. Prepare your computer with a static IP address:
   - For access over an Ethernet network port, use a static IP address in the 169.254.0.0 subnet with subnet mask 255.255.0.0. For example, use 169.254.100.201 for your computer.
   - For access over the OOB port, use a static IP address in the 192.168.0.0 subnet with subnet mask 255.255.0.0. For example, use 192.168.0.201 for your computer.
2. Connect an Ethernet cable from an Ethernet port on your computer to either an Ethernet network port on the switch or to the OOB port on the switch.
3. If you are using the OOB port, reboot the switch so that the IP address for the OOB port is set to 192.168.0.239 with 255.255.255.0 as the subnet mask.
4. Launch a web browser such as Google Chrome, Mozilla Firefox, or Microsoft Edge.

5. Enter the default IP address of the switch in the web browser address field:
   - For access over an Ethernet network port, enter **169.254.100.100**.
   - For access over the OOB port, enter **192.168.0.239**. The login page displays.

   NOTE: The default web browser protocol is HTTPS, but you can still access the login page using HTTP.

6. Click the **Main UI Login** button.
   The main UI login page displays in a new tab.
7. Enter **admin** as the user name, enter your local device password, and click the **Login** button.
   The first time that you log in, no password is required. However, you then must specify a local device password to use each subsequent time that you log in.
   The System Information page displays.

Log in to the main UI with a known IP address

If you did not assign a static IP address to the switch but let a DHCP server in your network assign an IP address to switch, determine the IP address by accessing the DHCP server or by using an IP scanner utility.
The procedures in this manual assume that you know the IP address of your switch.

To use a known IP address to access the switch over the main UI:

1. Launch a web browser and enter the IP address of the switch in the address field of your web browser.
   If you do not know the IP address of the switch, see Log in to the main UI with a web browser on page 3 and the subsections.

   NOTE: The default web browser protocol is HTTPS. Do not use HTTP to access the UI.

   The login page displays.
2. Click the **Main UI Login** button.
   The main UI login page displays in a new tab.

3. Enter **admin** as the user name and enter one of the following passwords:
   - Enter your local device password. The first time that you log in, no password is required. However, you then must specify a local device password to use each subsequent time that you log in.
   - If you are managing the switch through the Engage Controller, enter the site password for the Engage site to which the switch is onboarded.
   - If, in addition to the main UI, you are also managing the switch through the Insight Cloud Portal or Insight app, enter the Insight network password for the Insight network to which the switch is added.

   For more information about the password, see Credentials for the main UI on page 13.
4. Click the **Login** button.
   The System Information page displays.

Main UI buttons and user-defined fields

The following table shows the command buttons that are used on the pages in the main UI:

Table 1.Main UI command buttons

Button	Function
Add	Clicking the Add button adds the new item configured in the heading row of a table.
Apply	Clicking the Apply button sends the updated configuration to the switch. Configuration changes take effect immediately.
Cancel	Clicking the Cancel button cancels the configuration on the page and resets the data on the page to the previous values of the switch.
Delete	Clicking the Delete button removes the selected item.
Refresh	Clicking the Refresh button refreshes the page with the latest information from the device.
Save	Clicking the **Save** button saves your settings. This is located in the upper right corner of the device UI.
Logout	Clicking the **Logout** button ends the session. This is located in the upper right corner of the device UI.

User-defined fields can contain 1 to 159 characters, unless otherwise noted on the configuration web page. All characters can be used except for the following (unless specifically noted in for that feature):

Table 2.Invalid characters for user-defined fields

Invalid characters for user-defined fields						
\	|	/	<	>	*	?

Interface naming conventions

The switch supports physical and logical interfaces. Interfaces are identified by their type and the interface number. The physical ports are multispeed Ethernet interfaces or fiber ports and are numbered on the front panel. You configure the logical interfaces.
The following table describes the naming convention for all interfaces available on the switch.

Table 3.Naming conventions for interfaces

Interface	Description	Example
Physical interfaces	Depending on the switch model, the physical ports are multispeed 10G Ethernet ports, multispeed 2.5G Ethernet ports, Gigabit Ethernet ports, 25G fiber ports, 10G fiber ports, or a combination of such ports. The interface number consists of the switch unit number (from 1 to 8) followed by a forward slash, the slot number (which is always 0) followed by a forward slash, and the port number, which is a sequential number starting from 1.	1/0/1, 1/0/2, 1/0/3, and so on For a switch stack: 2/0/1, 2/0/2, 2/0/3, and so on 3/0/1, 3/0/2, 3/0/3, and so on
Link aggregation group (LAG)	A LAG interface is a logical interface that aggregates two or more physical ports.	LAG 1, LAG 2, LAG 3, and so on
CPU management interface	This is the internal switch interface responsible for the switch's base MAC address. This interface is not configurable and is always listed in the MAC Address Table.	0/5/1
Routing VLAN interfaces	This is an interface used for routing functionality.	VLAN 1, VLAN 2, VLAN 3, and so on

Save your settings to the running configuration

When you click the **Apply** button, your changes are saved for the web management session but

are not retained when you restart the switch. To save your changes to the running configuration (that is, permanently), click the **Save** icon at the top right of a page.
You can also first make multiple changes without clicking the **Save** icon after each change (although you must click the **Apply** button after each change) and then save the configuration to the running configuration. For more information, see Save the configuration on page 762).

Main UI online help

When you log in to the switch, each page contains a link to the online help that contains information to assist in configuring and managing the switch. The online help pop-up windows are context sensitive. For example, if the IP Addressing page is open, the help topic for that page displays if you click the Help button.
You can connect to the online support site at netgear.com/support/enterprise when you are logged in to the switch.

To access the online support link:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Help > Online Help > Support**.
3. To connect to the NETGEAR support site for the M4250 Series switches, click the **Apply** button.

Use the Device View in the Main UI

The Device View is an HTML applet that displays the ports on the switch. This graphic provides an alternate way to navigate to configuration and monitoring options. The graphic also provides information about device ports, current configuration and status, tables, and feature components.

To use the Device View:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **System > Device View**.
   The Device View front view of the switch displays. If a stack is configured, the front view of each stacked switch displays.
   The port coloring indicates the port status:
   - **Black**: The port is disabled or not connected to a device.
   - **Green**: The port is connected to a device.
   - **Gray**: The port is detached. For example, a port is detached if it is configured as a stack port but not connected to a stack member.
   - **Red**: An error occurred on the port.

   The port LEDs indicate the port status and the system LEDs indicate the system status. For more information about the LEDs, see the hardware installation guide, which you can download by visiting netgear.com/support/download/.
3. To display the back view, click the **B** button, and to display the front view again, click the **F** button.
   If the front view displays, the B button displays so that you can switch to the back view. If the back view displays, the F button displays so that you can switch to the front view.
4. To display the menus that let you configure ports, point to a port and right-click the port.
   Using the menu in the Device View, you can navigate to a page and configure the port settings.
5. To display the menus that let you configure global switch settings, point to anywhere on the switch other than a port, and right-click.
   Using the menu in the Device View, you can navigate to a page and configure the switch settings.

Set up SNMPv3 access

The switch supports the configuration of SNMP groups and users that can manage traps that the SNMP agent generates.
The switch uses both standard public MIBs for standard functionality and private MIBs that support additional switch functionality. All private MIBs begin with a "-" prefix. The main object for interface configuration is in -SWITCHING-MIB, which is a private MIB. Some interface configurations also involve objects in the public MIB, IF-MIB.
SNMP is enabled by default. The System Information page, which is the page that displays when you log in, displays the information that you need to configure an SNMP manager to access the switch.
You cannot access the switch using SNMPv3 until you log in to the switch as an admin and change the default password (see Log in to the main UI with a web browser on page 3). After you do, SNMPv3 is automatically configured with the SHA512 authentication protocol and the new password for admin user.
For SNMPv3 switch access, the authentication protocol must be SHA512.
For more information about SNMP, see Simple Network Management Protocol on page 158.

To configure authentication and encryption settings for the SNMPv3 admin profile:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > SNMP > SNMP V3 > User Configuration**.
   The User Configuration page displays.
3. From the **Authentication Protocol** menu, select **SHA512** for the authentication.
   SHA512 is the only possible choice from the menu.
4. In the **Authentication Key** field, type a password (key).
   The password can be up to 32 characters.
5. From the **Encryption Protocol** menu, select **None** or **AES128** for the encryption:
   - **None**: The information is not encrypted.
   - **AES128**: You must specify an encryption password for SNMPv3 access (see the following step).
6. If you select **AES128** from the **Encryption Protocol** menu, specify a password (key) in the **Encryption Key** field.
   The password can be up to 32 characters.
7. Click the **Apply** button.
   Your settings are saved.

8. To save the settings to the running configuration, click the **Save** icon.

NOTE: For information about using SNMPv1 or SNMPv2, see Simple Network Management Protocol on page 158.

Discover and onboard the switch to the NETGEAR Engage Controller

You can use the NETGEAR Engage Controller application to discover, onboard, manage, and monitor the switch. The NETGEAR Engage Controller application (referred to as the *controller* in this section) provides central management and configuration of switches and other supported devices through an audio-video (AV)-friendly, portable app for Windows and MacOS.

NOTE: This section refers to a setup where the controller is already installed. To download and install the controller, see the information in the *Engage Controller User Manual*, which you can download by visiting the NETGEAR Download Center at netgear.com/support/enterprise.

To discover and onboard the switch to the controller:

1. On your computer, in the folder in which you installed the controller application, double-click the **Engage** application icon, or double-click the **Engage** shortcut.
   The login window opens.
2. In the **Login Name** field, enter **admin** as the user name.
3. In the **Password** field, enter the controller password that you set up the first time that you logged in.
4. Click the **Login** button.
   The Managed Devices page displays.
5. If you set up more than one site, do the following:
   a. From the **Site** menu, select the site.
   b. In the Edit Network Setup pop-up window, confirm that the IP address settings and Engage network interface for the site are still correct by clicking the **Apply** button. You can also change the IP address settings and Engage network interface for the site.

6. Depending on whether the switch displays in the Discovered Devices table, do one of the following:
   - **The switch displays in the Discovered Devices table**: Click the **Onboard** button for the switch to display the Onboard pop-up window.
   - **The switch does not display in the Discovered Devices table**: Click the **Add Device** link to display the Add Device pop-up window.
7. To let the controller onboard the switch so that it can control the switch, specify the following information:
   - **Device Password**: Enter the local device password for the switch. If the switch is in the factory default state, click the **Use device default password** toggle so that it display green and is positioned to the right.
   - **Device IP Address**: If the IP address is not listed, type the IP address for the switch.

   NOTE: By default, the Device Admin User field displays *admin*.
8. Depending on the pop-up window that you are using, do one of the following:
   - On the Onboard pop-up window, click the **Apply** button.
   - On the Add Devices pop-up window, click the **Add** button.

   Your settings are saved.
9. To save the settings to the running configuration, at the top right of the page, click the **Save** button.

   You can now manage and monitor the switch by using the controller. For more information, see the *Engage Controller User Manual*.

NOTE: After you onboard the switch to an Engage site, the Engage site password replaces the switch local device password for the main UI. To access the main UI, you must then enter the Engage site password.

How Insight and the main UI interact with each other

If you manage the switch using the NETGEAR Insight Cloud Portal or Insight app, you can *also* still use the main UI to manage the switch. That is, these management methods are not mutually exclusive but complement each other. Changes to Insight are synchronized to the main UI, and the other way around, changes to the main UI are synchronized to Insight.

NOTE: Synchronization between the main UI and Insight might take up to 15 minutes. During the synchronization period, do not make the same changes on both the main UI and Insight or conflicting changes on the main UI and Insight.

By default, NETGEAR Insight is enabled but inactive if you do not configure and use it. If Insight is inactive, changes to the main UI are not synchronized to Insight. You cannot use the main UI to disable Insight. However, if you do not configure and use Insight, it is effectively disabled.

NOTE: NETGEAR Insight and the Engage Controller application are mutually exclusive.

Limitations

NETGEAR Insight does not support the following features, which you can configure in the main UI:

- SNTP servers
- IP routing mode
- MLD VLAN configuration

In addition, NETGEAR Insight does not support specific settings for the following features, which are fully configurable in the main UI:

- PoE global settings are supported, except for the system usage threshold and the power management mode
- PoE port settings are supported, except for port priority
- Voice VLAN is supported, except for port interface mode and port value
- Auto-VoIP is supported, but only the default OUI table (no custom OUIs)
- STP is supported, but not MSTP
- VLAN routing is supported, but not a configurable IP MTU
- Route configuration is supported, except for the route preference and route description
- CoS interface configuration is supported, except for the interface shaping rate
- RADIUS servers are supported, except for the RADIUS accounting server configuration
- Traffic control is supported, except for the MAC filter configuration and port security settings
- ACL configurations are supported, but the ACL Wizard is not
- HTTP file export is supported, but not for a text configuration; HTTP file upgrade is supported, but not for a text configuration

You can configure these features in the main UI.

Credentials for the main UI

The information in this section refers to the local device password that you must use when you access the switch main UI.

The first time that you access the main UI, enter the default switch admin password (password), after which you are required to customize the password for greater security. Subsequent times that you log in to the main UI, use your customized switch admin password.

If you onboard the switch to the NETGEAR Engage Controller application and then add the switch to an Engage controller site, the local device password changes to the Engage controller *site password*.

If you add the switch to a NETGEAR Insight network, the local device password changes to the

Insight *network password*.
For information about how the Insight network location password functions and for knowledge base articles about NETGEAR Insight, visit https://kb.netgear.com/000044336.
The following table lists the credential options for access to the main UI in relation to the Engage controller.

Table 4.Credentials for access to the main UI in relation to the Engage controller

Onboarded to an Engage controller site	Credentials	Main UI menu
No	Local device password	Full main UI menu
Yes	Engage controller site password	Full main UI menu
First added to a site and then removed from both the site and the controller	Engage controller site password until you set a new local device password	Full main UI menu

The following table lists the credential options for access to the main UI in relation to Insight.

Table 5.Credentials for access to the main UI in relation to Insight

Added to an Insight network	Credentials	Main UI menu
No	Local device password	Full main UI menu
Yes	Insight network password	Full main UI menu
First added to an Insight network, then removed from the network but not added to another Insight network.	Insight network password until you set a new local device password	Full main UI menu

NOTE: NETGEAR Insight and the Engage Controller application are mutually exclusive.

Manage Stacking

This chapter covers the following topics:

- Switch stacking overview
- Stack firmware synchronization and upgrade
- Stack configuration maintenance
- Stack management switch election
- Stack factory default reset behavior
- Stacking limitations
- Configure a stack
- Stack nonstop forwarding

Switch stacking overview

A stackable switch is a switch that is fully functional operating as a stand-alone unit but can also operate together with up to sevenfive other switches. This group of switches has the characteristics of a single switch and the port capacity of the sum of all stacked switches.
One of the switches in the stack controls the operation of the stack. This switch is called the stack management switch. The other switches in the stack are stack members. The stack management switch and stack members use stacking technology to behave and work together as a unified system. Layer 2 and higher protocols present the entire switch stack as a single entity to the network.
A switch stack can be described in terms of three semi-independent functions:

- **Forwarding plane**: The forwarding plane forwards data packets. The forwarding plane is implemented in hardware.
- **Control plane**: The control plane is the set of protocols that determine how the forwarding plane forwards packets, that is, which packets are allowed to be forwarded and where they go. Application software on the management switch acts as the control plane.
- **Management plane**: The management plane is application software running on the management switch that provides the interfaces, allowing you to configure and monitor the switch stack.

The stack management switch is the single point of stack-wide management. From the stack management switch, you configure the following:

- System-level (global) features that apply to all stack members
- Interface-level features for all interfaces on any stack member

A switch stack is identified in the network by its network IP address. The network IP address is assigned according to the MAC address of the stack management switch. Every stack member is uniquely identified by its own stack member number, which is a number from 1 to 61 to 8. The stack management switch can be any number within that range.
Stacking supports the following:

- Up to eight switches per stack.
- Up to six switches per stack.
- Single IP address management through a web browser, the CLI, or SNMP.
- Management switch-member configuration:
  - The stack management switch retains the configuration for the entire stack.
  - Automatic detection of new members, with synchronization of firmware (upgrade or downgrade, as needed).
- Configuration updates across the stack through a single operation.
- Automatic management switch failover. Fully resilient stack with ring topology or spine and leaf topology.
- Hot swapping (insertion and removal) of stack members.

Stack firmware synchronization and upgrade

All stack members must run the same firmware version to ensure compatibility within the stack. By default, if a switch is added to the stack and its firmware version is different from the firmware version on the stack management switch, that switch cannot join the stack. You can enable the Stack Firmware Auto Upgrade feature to automatically synchronize the firmware version on the new switch with the version running on the stack management switch. The synchronization operation might result in either upgrade or downgrade of firmware on the mismatched stack member.
Upgrading the firmware on a stack of switches is the same as upgrading the firmware on a single switch. After you download the firmware, see Update software or download a file on page 771, the downloaded firmware is distributed to all the connected switches of the stack. If you are setting up a new stack, we recommend that you first upgrade the firmware on the stack management switch and then add the stack members.

NOTE: We recommend that you set the active firmware for all stack members the same as the active firmware of the stack management switch. In other words, if image 1 is the active image on the stack management switch, all switches must use image 1 as the active image. For information about configuring the active image, see Change the software image that loads when the switch starts or reboots on page 781.

Stack configuration maintenance

The stack management switch stores and maintains the saved and running configuration files for the switch stack. The configuration files include the system-level settings for the switch stack and the interface-level settings for all stack members. Each stack member retains a copy of the saved file for backup purposes. If the management switch is removed from the stack or becomes unavailable, another member is elected management switch and then runs from that saved

configuration.
The stack management switch copies its running configuration to the stack member configured as the standby unit whenever it changes (subject to some restrictions to reduce overhead). This enables the standby unit to take over the stack operation with minimal interruption if the stack management switch becomes unavailable. The running-config synchronization also occurs when the running configuration is automatically saved on the stack management switch or when the standby unit changes.

Stack management switch election

All stack members are eligible stack management switches. If the stack management switch becomes unavailable, the remaining stack members participate in electing a new stack management switch from among themselves. The following factors determine which switch is elected as the stack management switch:

- Current role: The switch that is management switch always has priority to retain the role of management switch.
- Assigned priority
- MAC address

When the stack is powered up and completes the boot process or the original stack management switch becomes unavailable, the stack management switch is determined through an election process.
The rules for stack management switch election are as follows:

- If a unit was elected stack management switch previously, then it remains the stack management switch and other units are stack members.
- If no units were stack management switches, or more than one unit was a stack management switch, then the unit with the highest management priority is elected stack management switch. You can assign the management priority. However, if all units are assigned the same management priority, then the unit with the highest MAC address is assigned as the stack management switch.

Stack factory default reset behavior

If the stack management switch is reset to the factory default settings, see Reset the switch to the factory default settings on page 767, the stack management switch applies the default settings to all the stack members and resets the stack, including all participating stack members. When the stack starts up, the stack management switch election process begins.

Stacking limitations

A stacking link works only on the highest speed supported by a stack port. For example, a 25G port that is also capable of operating at 10G speed operates only at 25G speed in stack mode.

When you set up a stacking link between two switches, you must ensure that both ports operate at the same speed. For example, set up a stacking link between two ports that each support a maximum speed of 25G, or between two ports that each support a maximum speed of 10G.

NOTE: In Ethernet mode, this limitation does not apply, and the port can operate at 25G or a lower speed.

The following features are not compatible with stacking, so you cannot enable these features on a stack:

- Precision Time Protocol (see Precision Time Protocol transparent clock on page 75 and Precision Time Protocol boundary clock and grandmaster clock on page 78)
- 802.1AS audio video bridging (see 802.1AS timing and synchronization on page 323)
- Multiple Registration Protocol and 802.1Qav (see Multiple Registration Protocol and 802.1Qav on page 329)

Configure a stack

You can configure a stack by configuring the management switch and adding a standby management switch and regular members to the stack. You can also set the stack sample mode, which defines how the timestamps of the heartbeat messages between the stack members are computed, and the stack trunk hash mode, which defines how messages between the stack members are secured.

NOTE: You also can reassign the stack management switch functionality from one switch to another. If you do so, the entire stack configuration (including all interfaces in the stack) is removed and reconfigured on the new management switch. After the reload is complete, all stack management is performed on the new management switch. To preserve the current configuration when reassigning stack management to a different switch, save the current configuration to the running configuration *before* you reassign the stack management to a different switch. A stack move causes all routes and Layer 2 addresses to be lost (they can be relearned in the new stack configuration).

Set the stack sample mode

The stack sample mode defines how the timestamps of the heartbeat messages between the stack members are computed. These heartbeat messages allow all members of the stack to remain synchronized. The stack sample mode and, if configured, sample size apply to all members in the stack.

To set the stack sample mode:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Stacking > Basic > Stack Configuration**.
   The Stack Configuration page displays. The page shows different sections.
3. In the Stack Sample Mode section, from the **Sample Mode** menu, select the mode:
   - **Cumulative**: Tracks the sum of received timestamp offsets cumulatively.
   - **History**: Tracks the history of received timestamps.
4. If you select **History** from the **Sample Mode** menu, in the **Max samples** field, type the maximum number of samples to keep.
   The range is from 100 to 500 samples.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Add a switch as the stack management switch

When you configure a stack, one switch must function as the stack management switch. You can configure more than one management switch. In that situation, the switch with the highest priority functions as the management switch.
For information about configuring other switch roles in a stack, see Add a stack member or standby management switch on page 21.
If you are changing the role of an existing stack member to the role of management switch, see Change the setting for a switch in the stack on page 22.

To add a switch as the stack management switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Stacking > Basic > Stack Configuration**.
   The Stack Configuration page displays. The page shows different sections.
3. From the **Unit ID** menu, select an ID from **1** to **8**.
   This is the ID for the management switch.
4. From the **Switch Type** menu, select the switch model.

5. Optionally, from the **Switch Priority** menu, select the priority of the switch that determines if and when the switch becomes the stack management switch in the automatic stack management election process.
   The priority is in relation to the priority setting for other switches in the stack. A higher value indicates a higher priority. The maximum value is 15. The default is unassigned. If you set the priority to zero, the switch cannot become a stack management switch. However, if the switch was assigned or elected stack management switch previously, then it remains the stack management switch and other units are stack members.
6. From the **Management Status** menu, select **Management**.
7. Click the **Add** button.
   Your settings are saved and the switch is added to the stack as the management switch.
8. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields in the section. The Standby Status field does not apply to the management switch.

Field	Description
Hardware Management Preference	The hardware management preference of the switch, which can be Disabled or Unassigned.
Standby Status	The standby management switch status of the switch: • **Cfg Standby**: The switch is configured as the standby management switch. If the current stack management switch fails, this switch becomes the stack management. • **Opr Standby**: The switch is operating as the standby management switch because the configured standby management switch failed or is no longer a member of the stack. • **None**: The switch is not configured as the standby management switch.
Switch Status	The status of the switch in the stack: • **OK**: The switch is connected and working correctly. • **Unsupported**: The switch is not supported. • **Code Mismatch**: The firmware version is different from the version on the stack management switch. • **Config Mismatch**: The switch mismatches the switches in the stack. • **Not Present**: The switch is not connected. • **SDM Mismatch**: The Switch Database Management (SDM) template does not match. • **Updating Code**: A firmware update is in progress. • **STM Mismatch**: A stack manager (STM) mismatch occurred.

Add a stack member or standby management switch

You can add a switch to a stack and configure the switch role as stack member or standby switch:

- **Stack member**: The switch is a stack member but not the management switch or a standby switch.
- **Standby management switch**: The standby management switch takes over as the stack management switch if the stack management switch becomes unavailable.

For information about adding the management switch, see Add a switch as the stack management switch on page 19.
For information about changing the settings of a switch in the stack, see Change the setting for a switch in the stack on page 22.

To add a stack member or standby management switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Stacking > Basic > Stack Configuration**.
   The Stack Configuration page displays. The page shows different sections.
3. Go to the Stack Configuration section.
   If you already added one or more switches as members to the stack, the members show in the table.
4. From the **Unit ID** menu, select the ID of the member.
5. From the **Switch Type** menu, select the switch model.
6. Optionally, from the **Switch Priority** menu, select the priority of the switch that determines if and when the switch becomes the stack management switch in the automatic stack management election process.
   The priority is in relation to the priority setting for other switches in the stack. A higher value indicates a higher priority. The maximum value is 15. The default is unassigned. If you set the priority to zero, the switch cannot become a stack management switch. However, if the switch was assigned or elected stack management switch previously, then it remains the stack management switch and other units are stack members.
7. From the **Management Status** menu, select if the switch must function as a stack member (**Stack Member**) or as the standby management switch (**Standby**).
8. Click the **Add** button.
   Your settings are saved. The switch is added as a stack member or as a standby switch.
9. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields in the section.

Field	Description
Hardware Management Preference	The hardware management preference of the switch, which can be Disabled or Unassigned.
Standby Status	The standby management switch status of the switch: • **Cfg Standby**: The switch is configured as the standby management switch. If the current stack management switch fails, this switch takes over stack management. • **Opr Standby**: The switch is operating as the standby management switch because the configured standby management switch failed or is no longer a member of the stack. • **None**: The switch is not configured as the standby management switch.
Switch Status	The status of the switch in the stack: • **OK**: The switch is connected and working correctly. • **Unsupported**: The switch is not supported. • **Code Mismatch**: The firmware version is different from the version on the stack management switch. • **Config Mismatch**: The switch mismatches the switches in the stack. • **Not Present**: The switch is not connected. • **SDM Mismatch**: The Switch Database Management (SDM) template does not match. • **Updating Code**: A firmware update is in progress. • **STM Mismatch**: A stack manager (STM) mismatch occurred.

Change the setting for a switch in the stack

You can change the settings for a switch in the stack, including the unit ID, priority in relation to other switches in the stack, and the role for the switch in the stack.

To change the settings for a switch in the stack:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Stacking > Basic > Stack Configuration**.
   The Stack Configuration page displays. The page shows different sections.
3. In the Stack Configuration section, select the check box for the unit ID of the stack member.

4. To set another ID for the member, in the **Change Switch ID to** field, type an ID.

   CAUTION: The renumbering process causes the switch to restart after you click the Apply button.

5. If you previously assigned an incorrect switch model, from the **Switch Type** menu, select the correct switch model.
6. To change the priority of the switch in relation to other switches in the stack, from the **Switch Priority** menu, select the priority of the switch that determines if and when the switch becomes the stack management switch.
   The priority is in relation to the priority setting for other switches in the stack. A higher value indicates a higher priority. The maximum value is 15. The default is unassigned. If you set the priority to zero, the switch cannot become a stack management switch. However, if the switch was assigned or elected stack management switch previously, then it remains the stack management switch and other units are stack members.
7. To change the role of the switch, from the **Management Status** menu, select if the switch must function as a normal stack member, the stack management switch, or the standby management switch:
   - **Stack Member**: The switch is a stack member but not the management switch or a the standby management switch.
   - **Management**: The switch is the stack management switch.
   - **Standby**: The switch is a standby management switch that takes over as the stack management switch if the stack management switch becomes unavailable.
8. Click the **Apply** button.
   Your settings are saved.
   If you set another unit ID for the switch or change the role of the switch in the stack, the switch reboots.
9. To save the settings to the running configuration, click the **Save** icon.

Remove a switch from the stack

You can remove a switch that is no longer required in the stack.

To remove a switch from the stack:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **System > Stacking > Basic > Stack Configuration**.
   The Stack Configuration page displays. The page shows different sections.
3. In the Stack Configuration section, select the check box for the unit ID of the stack member.
4. Click the **Delete** button.
   Your settings are saved and the switch is removed from the stack.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the stack mode for a port

By default, the Ethernet ports on a switch are configured as standard Ethernet ports for Ethernet traffic. However, you can configure an Ethernet port as a stack port that can connect to a stack port on another stack member.

To configure the stack mode for a port:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Stacking > Advanced > Stack-port Configuration**.
   The Stack-port Configuration page displays.
3. To display the ports of a specific stacked switch, select the unit ID for the stacked switch, or select the **All** link to display all ports for all stacked switches.
   The ports for the selected switch or switches display.
4. Select the check boxes for one or more ports.
5. From the **Configured Stack Mode** menu, select the port operating mode:
   - **Stack**: The port can connect to the stack port on another stacked switch.
   - **Ethernet**: The port operates as a standard Ethernet switch port that receives and transmits network traffic. This is the default setting.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
Unit ID	The unit ID for the stacked switch
Port	The port on the stacked switch
Type	The type of port
Product Name	The name of the fiber transceiver module installed, if any
Running Stack mode	The operational mode of the port. For example, you can configure a port for stacking, but until you restart the switch, the operational mode of the port is still Ethernet mode. After you restart the switch, the operational mode transitions into Stack mode.
Link Status	The link status of the port: Up or Down
Link Speed (Gbps)	The maximum speed of the port in Gbps
Transmit Data Rate (Mbps)	The approximate transmit rate in Mbps on the port
Transmit Error Rate (Errors/s)	The number of errors in transmitted packets per second
Total Transmit Errors	The total number of errors in transmitted packets since the switch restarted. The counter might restart.
Receive Data Rate (Mbps)	The approximate receive rate in Mbps on the port
Receive Error Rate (Errors/s)	The number of errors in received packets per second.
Total Receive Errors	The total number of errors in received packets since the switch restarted. The counter might restart.
Link Flaps	The total number of link flaps

Configure stack firmware synchronization

You can configure if stack firmware is automatically synchronized, if a firmware downgrade to an earlier version is allowed, and if SNMP traps are generated.

To configure stack firmware synchronization:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Stacking > Advanced > Stack Firmware Synchronization**.
   The Stack Firmware Synchronization page displays.
3. Select a Stack Firmware Auto Upgrade radio button to specify what happens if you add a new stack member and its firmware version does not match the version on the stack management switch:
   - **Enable**: The stack management switch updates the firmware version on the new stack member so that it matches the firmware version on the stack.
   - **Disable**: The stack management switch does not update the firmware version on the new stack member, which effectively prevents the new stack member from joining the stack.
4. Select a Traps radio button to specify what happens if you add a new stack member:
   - **Enable**: SNMP traps are sent during the start and completion of a stack firmware synchronization or after synchronization fails.
   - **Disable**: No SNMP traps are sent during a stack firmware synchronization.
5. Select an Allow Downgrade radio button to specify if a firmware version downgrade is allowed if you add a new stack member that runs a firmware version that is later than the firmware version of the stack:
   - **Enable**: The stack management switch downgrades the firmware version on the new stack member so that it matches the firmware version on the stack.
   - **Disable**: The stack management switch does not downgrade the firmware version on the new stack member, which can effectively prevent a switch from joining the stack.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Configure the trunk hash mode

The trunk hash mode determines the load-balancing mode that is used for a stack. After you select the trunk hash mode, the switch balances traffic by selecting one of the links over which specific packets must be transmitted. The switch selects the link by creating a binary pattern from selected fields in a packet and associating that pattern with a particular link.

To configure the trunk hash mode:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Stacking > Advanced > Stack-port Configuration**.
   The Stack-port Configuration page displays.
3. To display the ports of a specific stacked switch, select the unit ID for the stacked switch, or select the **All** link to display all ports for all stacked switches.
   The ports for the selected switch or switches display.
4. Select the check box for one or more ports:
5. From the **Stack Trunk Hash Mode** menu, select the mode:
   - **1 Src MAC, VLAN, EType, incoming port**: The source MAC address, VLAN, EtherType, and incoming port associated with the packet
   - **2 Dest MAC, VLAN, EType, incoming port**: The destination MAC address, VLAN, EtherType, and incoming port associated with the packet
   - **3 Src/Dest MAC, VLAN, EType, incoming port**: The source and destination MAC addresses, VLAN, EtherType, and incoming port associated with the packet
   - **4 Src IP and Src TCP/UDP Port fields**: The source IP address and source TCP/UDP port fields of the packet
   - **5 Dest IP and Dest TCP/UDP Port fields**: The destination IP address and destination TCP/UDP port fields of the packet
   - **6 Src/Dest IP and TCP/UDP Port fields**: The source and destination IP addresses and source and destination TCP/UDP port fields of the packet.
   - **7 Enhanced hashing mode**: Dynamic selections of fields based on the packet flow. For L2 packets, the source and destination MAC addresses are used; for IP packets, the source and destination IP addresses and TCP/UDP port fields are used.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Display information about the stack configuration

You can display information about the stack configuration.

To display information about the stack configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **System > Stacking > Basic > Stack Configuration**.
   The Stack Configuration page displays. The page shows different sections.
3. Go to the Basic Stack Status section.
   The following table describes the view-only fields in the section:

Field	Description
Unit ID	The unit ID of the switch in the stack
Switch Description	The description for the switch, which effectively is the switch model number
Serial Number	The serial number of the switch
Uptime	The period since the switch was restarted
Preconfigured Model Identifier	The model type assigned by the switch manufacturer, which effectively is the switch model number
Plugged-In Model Identifier	The model type assigned by the switch manufacturer to identify the plugged-in switch. (The Plugged-In Model Identifier can be the same as the Preconfigured Model Identifier.)
Detected Code Version	The detected version of the firmware of the switch
Detected Code in Flash	The detected version of the firmware of the switch in flash memory
SFS Last Attempt Status	The status of the most recent stack firmware synchronization (SFS).

Display the heartbeat message stack status

You can display the heartbeat message stack status and clear the sampling information. A heartbeat message is a message that a stacked switch exchanges with its neighbors.

To display the heartbeat message stack status and clear the sampling information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Stacking > Advanced > Stack Status**.
   The Stack Status page displays.

3. To display the stack status for a specific stacked switch, select the unit ID for the stacked switch, or select the **All** link to display the stack status for all stacked switches.
   The following table describes the view-only fields in the section:

Field	Description
Unit ID	The unit ID of the switch in the stack
Neighbor Unit ID	The ID of the neighboring unit with which data is exchanged
Current	The current time of heartbeat message reception
Average	The average time between two consecutive incoming heartbeat messages
Min	The minimum time between two consecutive incoming heartbeat messages
Max	The maximum time between two consecutive incoming heartbeat messages
Dropped	The number of dropped or lost heartbeat messages

4. To refresh the page, click the **Refresh** button.
5. To clear the sampling information, do the following:
   a. From the **Clear Counters** menu, select a unit ID, or select **All**.
   b. Click the **Apply** button. The counters are cleared.
6. To save the settings to the running configuration, click the **Save** icon.

Display stack port diagnostic, status, and path information

Stack port diagnostics lets you view diagnostic, status, and path information about a stack port.

To run stack port diagnostics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Stacking > Advanced > Stack-port Diagnostics**.
   The Stack-port Diagnostics page displays.
3. In the Stack-port Diagnostics section, to display the diagnostic information for a specific stacked switch, select the unit ID for the stacked switch, or select the **All** link to display diagnostic information for all stacked switches.
   The ports for the selected switch or switches display.
   For each stack port, the Port Diagnostics Info field displays diagnostic and status information.

4. In the Stack-port packet-path section, to display information about the packet path between stack ports, select the unit ID for the stacked switch, or select the **All** link to display diagnostic information for all stacked switches.
   The Direction field displays the stack unit number from where the stack path is computed.
   The Packet-path field displays the stack port number (for example, 2/0/28) that is connected to the stack unit from where the stack path is computed.
5. To refresh the page, click the **Refresh** button.
6. To clear the statistics, click the **Clear** button.
7. To save the settings to the running configuration, click the **Save** icon.

Remove a stack

You can remove a stack by removing all switches from the stack.

To remove a stack:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Stacking > Basic > Stack Configuration**.
   The Stack Configuration page displays. The page shows different sections.
3. In the Stack Configuration section, select all check boxes for the units in the stack.
   You cannot remove the stack management switch.
4. Click the **Delete** button.
   Your settings are saved and the units are removed from the stack.
5. To save the settings to the running configuration, click the **Save** icon.

Stack nonstop forwarding

Nonstop forwarding (NSF) allows the forwarding plane of stack members to continue to forward packets while the control and management planes restart as a result of a power failure, hardware failure, or software fault on the management switch. If the management switch fails, traffic flows can continue to enter and exit the stack through physical ports on any member other than the management switch with less than one second of interruption.
You can also initiate a nonstop forwarding failover to a standby management switch (see Configure nonstop forwarding and display associated information on page 31).
To prepare the standby management switch, applications on the management switch continuously forward information to the standby unit. Changes to the running configuration are automatically copied to the standby management switch. MAC addresses stay the same across a

nonstop forwarding failover so that neighbors are not required to relearn them.
When a nonstop forwarding failover occurs, the control plane on the standby management switch applies the information that it received from the management switch. When the control plane is initializing, the stack cannot react to external changes, such as network topology changes. After the control plane is fully operational on the new management switch (which was previously the standby management switch), the control plane updates hardware changes as needed. Control plane failover time depends on the size of the stack, the complexity of the configuration, and the speed of the switch CPU.
The management plane restarts when a failover occurs. Management connections must be reestablished.
For NSF to be effective, adjacent networking devices must not reroute traffic around the restarting device. The switch uses three techniques to prevent traffic from being rerouted:

- A protocol can distribute a part of its control plane to stack members so that the protocol can give the appearance that it is still functional during the restart. Spanning tree and port channels use this technique.
- A protocol can enlist the cooperation of its neighbors through a technique known as graceful restart. OSPF (see Configure OSPF and OSPFv3 on page 406) uses graceful restart if you enable that option.
- A protocol can restart after the failover if neighbors react slowly enough that they cannot detect the outage. The IP multicast routing protocols are a good example of this behavior.

To take full advantage of nonstop forwarding, Layer 2 connections to neighbors must be through port channels that span two or more stack members, and Layer 3 routes must be equal-cost multipath (ECMP) routes with next hops through physical ports on two or more stack members. The hardware can quickly move traffic flows from port channel members or ECMP paths on a failed unit to a surviving unit.

Configure nonstop forwarding and display associated information

Nonstop forwarding (NSF) allows a standby management switch to take over if the stack management switch goes down or becomes unavailable.

To configure nonstop forwarding and display associated information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Stacking > Advanced > NSF Summary**.
   The NSF Summary page displays.

3. Select an Admin Status radio button:
   - **Disable**: NSF is disabled.
   - **Enable**: NSF is enabled. This is the default setting. When NSF is enabled, the stack automatically selects a standby management switch based on the stack member settings (see Add a stack member or standby management switch on page 21). The stack management switch copies stack information to the standby management switch.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.
6. To manually initiate an NSF switchover from the stack management switch to the standby management switch, click the **Initiate Failover** button.
   The standby management switch becomes the stack management switch.

The following table describes the view-only fields on the page.

Field	Description
NSF Summary section	
Operational Status	Indicates if NSF is enabled on the standby management switch
Last Startup Reason	The reason of the most recent NSF switchover: • **Power On**: The stack management switch restarted, for example because of a power cycle or the execution of a reload command, and the standby management switch took over. • **Cold Admin Move**: The stack management switch reset all information without actually restarting, the standby management switch took over and began from a preinitialized state, and no data is retained from before the switchover. • **Warm Admin Move**: You pressed the **Initiate Failover** button or a CLI command for the standby management switch to take over. • **Auto Warm**: The stack management switch restarted because of a failure and the standby management switch took over. • **Auto Cold**: The standby management switch took over and user data traffic was interrupted. This situation might be caused by multiple failures occurring close together.
Time Since Last Restart	The period since the current management switch became the active management switch. For the standby management switch, the value is set to 0.
Restart In Progress	Indicates if a switchover from the stack management switch to the standby management switch is in progress. A restart is not considered complete until all information is fully reconciled.
Warm Restart Ready	The stack management switch is synchronized with the standby management switch, and the stack management switch can be restarted, if needed.
Copy of Running Configuration to Backup Unit	
Status	The status of copying the running configuration to the standby management switch
Time Since Last Copy	The period since the running configuration was copied to the standby management switch
NSF Support on Unit	
Unit ID	The unit ID of the switch in the stack

Field	Description
NSF Support	Displays if the switch in the stack is capable of supporting NSF

Display the nonstop forwarding checkpoint statistics

The checkpoint statistics are the statistics for the automated messages that are sent between the management switch and the standby management switch.

To display the nonstop forwarding checkpoint statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Stacking > Basic > Checkpoint Statistics**.
   The Checkpoint Statistics page displays.
3. To refresh the page, click the **Refresh** button.
4. To clear the statistics, click the **Clear** button.
5. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Field	Description
Messages Checkpoint	The number of messages that were sent from the stack management switch to the standby management switch
Bytes Checkpointed	The amount of data in Bytes that was sent from the stack management switch to the standby management switch
Time Since Counters Cleared	The period since the counters were reset
Checkpoint Message Rate	The number of seconds or fractions of seconds between measurements
Last 10-second Message Rate	The number of messages that were sent in the last 10-second measurement interval
Highest 10-second Message Rate	The highest number of messages that were sent in a 10-second measurement interval

Configure Switch System Information

This chapter covers the following topics:

- Switch system and slot information
- Loopback interface
- IPv4 management interfaces and VLANs
- IPv6 management interfaces and VLANs
- Time and NTP settings
- Precision Time Protocol transparent clock
- Manage PTP multiple domains
- Domain Name System
- Switch database management template
- Green Ethernet settings
- Bonjour settings
- Enable or disable the TFTP server capability
- IPv4 DHCP server
- DHCP relay
- DHCP Layer 2 relay
- UDP relay
- DHCPv6 server
- DHCPv6 relay interface
- Power over Ethernet
- Timer schedules
- Simple Network Management Protocol
- Link Layer Discovery Protocol
- Link Layer Discovery Protocol for Media Endpoint Devices
- Link dependency
- Industry Standard Discovery Protocol
- Configure the RGB LED color settings

Switch system and slot information

You can configure the view and configure the switch system information.

View and configure switch system information

When you log in, the System Information page displays. You can configure and view general device information.

To view and define system information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Define the following fields:
   - **System Name**: Enter the name to identify this switch. You can use up to 255 alphanumeric characters. The factory default is blank.
   - **System Location**: Enter the location of this switch. You can use up to 255 alphanumeric characters. The factory default is blank.
   - **System Contact**: Enter the contact person for this switch. You can use up to 255 alphanumeric characters. The factory default is blank.
   - **Login Timeout:** Specify how many minutes of inactivity can occur on a serial port connection before the switch closes the connection. Enter a number between 0 and 160 minutes. The factory default is 5. Entering 0 disables the time-out.
   - **Management Source Interface**: Select the management interface to use as the source interface for SNMP trap, syslog, DNS, TACACS+, RADIUS, sFlow, and NTP features. By default, the following options display in the menu:
     - **None**: The primary IP address of the originating (outbound) interface is used as the source address.
     - **VLAN 1**: The primary IP address of VLAN 1 is used as the source address. This is the default selection.
     - **Service Port**: The management port IP address is used as the source address.

     Depending on the configuration of your switch, the following options can display:
     - **Another VLAN ID**: The primary IP address of a VLAN other than VLAN 1 is used as the source address.
     - **Routing interface**: The primary IP address of a routing interface is used as the source address.
     - **Routing VLAN**: The primary IP address of a VLAN routing interface is used as the source address.
     - **Routing loopback interface**: The primary IP address of a routing loopback interface is used as the source address.
     - **Different**: Applications for which the source interface is configured separately.
3. Click the **Apply** button.
   Your settings are saved.
4. To save the settings to the running configuration, click the **Save** icon.

The following table describes the status information in the Application Information and System Information sections on the page.

Table 6.Application Information and System Information

Field	Description
Application Information	
App Name	The name of the application that runs on the switch. Examples of applications that might be running are AVUI, ConfigAgent, and discAgent.
App Status	The status of the application
Version	The version of the application
System Information	
Product Name	The product name (model name) of this switch
IPv4 Management Address	The IPv4 address and mask assigned to the management VLAN interface
IPv6 Management Address	The IPv6 address and mask assigned to the management VLAN interface
IPv4 Management Interface	The IPv4 management VLAN ID of the switch. Click the displayed Management VLAN ID value to jump to the configuration page. See Configure an IPv4 management VLAN on page 51.
IPv6 Management Interface	The IPv6 management VLAN ID of the switch. Click the displayed Management VLAN ID value to jump to the configuration page. See Configure an IPv6 management VLAN on page 57.
IPv4 Loopback Interface	The IPv4 address and mask assigned to the loopback interface
IPv6 Loopback Interface	The IPv6 prefix and prefix length assigned to the loopback interface
System Date	The current date
System SNMP OID	The base object ID for the switch enterprise MIB
System MAC Address	Universally assigned MAC network address

View the fan status

This page shows the status of the fans in all units. These fans remove the heat generated during operation by PoE power ports (for models that support PoE), the CPU, and other internal components, to allow the chipsets to function normally.
The fan status can be one of the following:

- **Operational**: The fan is running normally.
- **Failure**: The fan failed.
- **Not Present**: The fan is absent. (The number of fans depends on the switch model.)
- **Stop**: The fan stopped because the operating temperature is low. The fan will start if the operating temperature rises above a preset threshold.

To view the fan status:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Scroll down to the Fan Status section.

3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fan status fields. The table can show the fan status for up to eight stacked switches. The number of fans depends on the switch model.

Table 7.Fan status information

Field	Description
Unit ID	The unit ID of the switch in which the fan is installed.
FAN	The fan number
Description	The fan description, such as FAN-1, FAN-2, and so on.
Type	The type of fan, which is always fixed.
Speed	The speed of the fan. For more information, see the hardware installation guide, which you can download by visiting netgear.com/support/download/.
Duty level	The duty level of the fan in percentage.
FAN Status	The status of the fan: ○ **Operational**: The fan is running normally. ○ **Failure**: The fan failed. ○ **Not Present**: The fan is not absent. (The number of fans depends on the switch model.) ○ **Stop**: The fan stopped because the operating temperature is low. The fan will start if the operating temperature rises above a preset threshold.

View the temperature sensor information

You can view the current temperature of different system sensors using the Temperature Status table.

To view temperature information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Scroll down to the Temperature Status section.

3. To refresh the page, click the **Refresh** button.
   The following table describes the temperature status view-only fields.

Table 8. Temperature status information

Field	Description
Unit ID	The unit ID of the switch in which the sensor is installed.
Sensor	The ID of the sensor.
Description	The description of the sensor.
Temperature (C)	The temperature of the sensor in Celsius.
State	The state of the sensor, which should be Normal.
Max Temp (C)	The maximum supported operating temperature in Celsius.

View the device status and firmware version

You can view the device status and firmware version.

To view the device status and firmware version:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Scroll down to the Temperature Status section.

3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only device status fields. The table can show the device status for up to eight stacked switches.

Table 9.Device status information

Field	Description
Unit ID	The unit ID of the switch for which the device status displays.
Firmware Version	The release, version, and maintenance number of the firmware running on the switch. For example, if the release is 1, the version is 2, and the maintenance number is 4, the format is 1.2.4.
Boot Version	The version of the boot code that is in the flash memory. The boot code enables the firmware to be loaded into the memory.
Serial Number	The serial number of the switch.
PS-1	The status of the internal power supply: ○ **Operational**: The power supply is present and functioning properly. ○ **Powering**: The main power failed or the power cable is disconnected but a redundant power supply provides power. ○ **Not Present**: No power supply is present in the slot. ○ **Not powered**: The power supply is present but not connected to a power source. ○ **Not powering**: The power supply is present and connected but the switch uses another power source. ○ **Incompatible**: The power supply is present but incompatible. ○ **Failed**: The power supply is present but the power cable is not plugged-in or a bad cable is plugged-in. ○ **N/A**: The switch does not support the power supply.
For PoE models only: MAX PoE	For PoE models only: The PoE system status on the switch: ○ **ON**: Less than 10W of PoE power is available for another PD. ○ **OFF**: At least 10W of PoE power available for another PD. ○ **N/A**: PoE is not supported by the switch.
System Up Time	The time in days, hours, and minutes since the switch was restarted.

View the system CPU status

To view the system CPU status:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > System CPU Status**.
   The System CPU Status page displays.
3. If you configured a switch stack, from the **Unit No** menu, select the unit ID of the switch in the stack.
   The CPU memory status includes the total memory of the switch in KBytes and the available memory space for the switch in KBytes.
   The CPU utilization shows the CPU utilization by the various processes on the switch.

Configure the CPU thresholds

You can configure CPU thresholds that trigger a notification if exceeded. The notification occurs through SNMP trap and syslog messages.

To configure the CPU thresholds:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Management>System CPU Status > CPU Threshold**.
   The CPU Threshold page displays.
3. Configure the following CPU threshold settings:
   - **Rising Threshold**: Configure the rising threshold value. A notification is generated when the total CPU utilization exceeds this threshold value over the configured period. The range is a percentage from 1 to 100.
   - **Rising Interval**: Configure the rising interval value. Configure this utilization monitoring period from 5 to 86400 seconds in multiples of 5 seconds.
   - **Falling Threshold**: Configure the falling threshold. A notification is triggered when the total CPU utilization falls below this level for a configured period of time. The falling utilization threshold must be equal to or less than the rising threshold value. The falling utilization threshold notification is generated only if a rising threshold notification was generated previously. Configuring the falling utilization threshold and period is optional. If the falling CPU utilization values are not configured, they take the same value as the rising CPU utilization values. The range is a percentage from 1 to 100.
   - **Falling Interval**: Configure the falling interval. You can configure the utilization monitoring period from 5 seconds to 86400 seconds in multiples of 5 seconds.
   - **Free Memory Threshold**: Configure the CPU free memory threshold value in KB.

4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

View or clear switch statistics

To view or clear the switch statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Switch Statistics**.
   The Switch Statistics page displays.
3. To clear all the counters, resetting all switch summary and detailed statistics to default values, click the Clear button.
   The discarded packets count cannot be cleared.
4. To save the settings to the running configuration, click the **Save** icon.

The following table describes switch statistics information.

Table 10.Switch statistics information

Field	Description
ifIndex	The ifIndex of the interface table entry associated with the processor of this switch.
Octets Received	The total number of octets of data received by the processor (excluding framing bits but including FCS octets).
Packets Received Without Errors	The total number of packets (including broadcast packets and multicast packets) received by the processor.
Unicast Packets Received	The number of subnetwork-unicast packets delivered to a higher-layer protocol.
Multicast Packets Received	The total number of packets received that were directed to a multicast address. This number does not include packets directed to the broadcast address.
Broadcast Packets Received	The total number of packets received that were directed to the broadcast address. This does not include multicast packets.
Receive Packets Discarded	The number of inbound packets that were discarded even though no errors were detected to prevent their being deliverable to a higher-layer protocol. A possible reason for discarding a packet could be to free up buffer space.
Octets Transmitted	The total number of octets transmitted out of the interface, including framing characters.
Packets Transmitted Without Errors	The total number of packets transmitted out of the interface.
Unicast Packets Transmitted	The total number of packets that higher-level protocols requested that are transmitted to a subnetwork-unicast address, including those that were discarded or not sent.
Multicast Packets Transmitted	The total number of packets that higher-level protocols requested that are transmitted to a multicast address, including those that were discarded or not sent.
Broadcast Packets Transmitted	The total number of packets that higher-level protocols requested that are transmitted to the broadcast address, including those that were discarded or not sent.
Transmit Packets Discarded	The number of outbound packets that were discarded even though no errors were detected to prevent their being

Field	Description
	deliverable to a higher-layer protocol. A possible reason for discarding a packet could be to free up buffer space.
Most Address Entries Ever Used	The highest number of Forwarding Database Address Table entries learned by this switch since the most recent reboot.
Address Entries in Use	The number of learned and static entries in the Forwarding Database Address Table for this switch.
Maximum VLAN Entries	The maximum number of virtual LANs (VLANs) allowed on this switch.
Most VLAN Entries Ever Used	The largest number of VLANs that were active on this switch since the last reboot.
Static VLAN Entries	The number of presently active VLAN entries on this switch that were created statically.
Dynamic VLAN Entries	The number of presently active VLAN entries on this switch that were created by GVRP registration.
VLAN Deletes	The number of VLANs on this switch that were created and then deleted since the last reboot.
Time Since Counters Last Cleared	The elapsed time, in days, hours, minutes, and seconds, since the statistics for this switch were last cleared.

View USB device information

To display the USB device information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > USB Device Information**.
   The USB Device Information page displays.
   The Device Status field displays the current status of the device. The status is one of the following:
   - **Active**: The device is plugged in and recognized by the switch.
   - **Inactive**: The device is not mounted.
   - **Invalid**: The device is not present or an invalid device is plugged in.

3. To refresh the page, click the **Refresh** button.
   The following table describes the information in the USB Memory Statistics section.

Table 11.USB Memory Statistics information

Field	Description
Total Size	The USB flash device storage size in bytes.
Bytes Used	The size of memory used on the USB flash device.
Bytes Free	The size of memory free on the USB flash device.

The following table describes the information in the USB Directory Details section.

Table 12.USB Directory Details information

Field	Description
File Name	The name of the file stored in the USB flash drive.
File Size	The size of the file stored in the USB flash drive in bytes
Modification Time	The last modification time of the file stored in the USB flash drive.

Display information about a stack and optional switch models for a stack

You can display information about an existing stack of M4350 series switches and optional switch models that are compatible with a stack of M4350 series switches. Although the page suggests that the information is for a slot configuration, it is really for a switch stack configuration.

To display information about a stack and optional switch models for a stack:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Slot Information**.
   The Slot Information page displays. The page shows different sections.
3. To refresh the page, click the **Refresh** button.

The following table describe the view-only fields in the Slot Configuration table.

Table 13.Slot configuration information (stacked switch configuration)

Field	Description
Slot	The unit and slot number. This entry represents a stacked switch. The first switch is 1/0, the second switch is 2/0, the third switch is 3/0, and so on.
Status	Indicates if the switch is configured for stacking (“Full”) or preconfigured for stacking (“Empty”)
Administrative State	The administrative state is always Enable
Power State	The power state is always Enable
Configured Card Model ID	The model of the M4350 series switch
Configured Card Description	The description of the M4350 series switch
Inserted Card Model ID	This field is identical to the Configured Card Model ID field
Inserted Card Description	This field is identical to the Inserted Card Description field
Card Power Down	Indicates if the power state can be administratively enabled or disabled. This field is always False.
Card Pluggable	Indicates if the switch is pluggable. This field is always False.

The following table describe the view-only fields in the Supported Card table.

Table 14.Supported card information (information about supported M4350 series switches for a stack)

Field	Description
Card Model	The model ID of the supported switch
Card Index	The index assigned to the supported switch
Card Type	The hardware type of the supported switch, which is assigned by the manufacturer
Card Descriptor	The description of the supported switch

The following table describe the view-only fields in the Supported Switch table. If you preconfigure a new stack member, the switch index identifies the type of switch that is being added to the stack.

Table 15.Supported switch information (switch index and management preference information)

Field	Description
Switch Model ID	The model number of the supported switch
Switch Index	The index that is assigned to the supported switch
Management Preference	The management preference of the supported switch

Loopback interface

A loopback interface lets you send a ping, traceroute, and other test traffic as well as protocol traffic to the switch. A loopback interfaces requires a unique IPv4 or IPv6 network address. You can add one or more loopback interfaces.

To add a loopback interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Loopback Interface**.
   The Loopback Interface page displays.
3. From the **Loopback Interface Type** menu, select IPv4 or IPv6.
4. From the **Loopback ID** menu, select an interface ID.
5. Configure the following settings, depending on the type of loopback interface:
   - **IPv4**: For an IPv4 interface, configure the following settings:
     - **Primary IP Address**: Enter the primary IPv4 address for this interface in dotted-decimal notation.
     - **Primary IP Subnet Mask**: Enter the primary IPv4 subnet mask in dotted-decimal notation.
   - **IPv6**: For an IPv6 interface, configure the following settings:
     - **IPv6 mode**: Select **Enable** to use the IPv6 address after you specify it. You can select this option only before you specify the IPv6 address.
     - **IPv6 Prefix**: Enter the IPv6 prefix.
     - **Prefix Length**: Enter the length for the IPv6 prefix.
     - **EUI64**: Enable or disable the 64-bit extended unique identifier (EUI-64).
6. Click the **Add** button.
   Your settings are saved and the loopback interface is added.
   The Loopback Interface Status fields show if the loopback interface is up.

7. To save the settings to the running configuration, click the **Save** icon.

IPv4 management interfaces and VLANs

The main UI provides IPv4 configuration options that you can use to set up either VLAN-based or port-based IP management.
By default, the source interface for applications is VLAN 1. If you change the IPv4 management VLAN and port values, the source interface to the VLAN 1 default VLAN and port also change automatically.
You can access the switch over the main UI by one of the following methods:

- **IPv4 service port**: The service port is a dedicated Ethernet port for out-of-band (OOB) management of the switch. Management traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.
- **Management VLAN interface**: The management VLAN is the logical interface used for in-band connectivity with the switch over any of the switch's network interfaces. To access the switch over a network you must first configure an IP address and subnet mask for the switch.
- **Optional single management interface**: By default, you can use any of the switch's network interfaces as the management interface. However, you can restrict IPv4 management to one specific network interface.

Configure the IPv4 service port

You can configure network information on the IPv4 service port, also referred to as the out-of-band (OOB) port. The service port is a dedicated Ethernet port for out-of-band management of the switch. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.
By default, no IP address is set for the OOB port, but its DHCP client is enabled so that the port can receive an IP address from a DHCP server in your network. If the OOB port does not receive an IP address from a DHCP server in your network, the IP address for the port is set to 192.168.0.239 with 255.255.255.0 as the subnet mask. The same occurs if you connect the OOB port directly to a computer and reboot the switch.

To configure the IPv4 service port:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Management>Management Interfaces > IPv4 Service Port Configuration**.
   The IPv4 Service Port Configuration page displays.

3. Select a Service Port Configuration Protocol radio button:
   - **Bootp**: During the next boot cycle, the BootP client on the device broadcasts a BootP request to acquire information from a BootP server on the network. The IP Address, Subnet Mask, and Default Gateway field information is dynamically acquired.
   - **DHCP**: During the next boot cycle, the DHCP client on the device broadcasts a DHCP request to acquire information from a DHCP server on the network. The IP Address, Subnet Mask, and Default Gateway field information is dynamically acquired.
   - **None**: The device does not attempt to acquire network information dynamically. You must configure the address settings manually:
     - **IP Address**: Specify the IP address of the interface.
     - **Subnet Mask**: Specify the IP subnet mask for the interface.
     - **Default Gateway**: Specify the default gateway for the IP interface.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 16.IPv4 service port configuration information

Field	Description
Burned-in MAC Address	The burned-in MAC address used for out-of-band connectivity.
Interface Status	Indicates whether the link status is up or down.
DHCP Client Identifier	If the selection is DHCP, the DHCP client ID.

Configure an IPv4 management VLAN

For you to manage the switch over the main UI, you must define the management IP address. A management VLAN interface is created by default and it gets an IP address if a DHCP server is present. If the management VLAN interface cannot get an IP address, the automatically assigned fallback IPv4 address is 169.254.100.100 and the subnet mask is 255.255.0.0.
A management VLAN is used as the default source interface for the syslog, the message log, an SNMP client, and so on. The network interface is disabled by default.
The management VLAN is the logical interface used for in-band connectivity with the switch through any of the switch's front panel ports. The configuration parameters associated with the switch's management VLAN do not affect the configuration of the front panel ports through which traffic is switched or routed.
To access the switch over a network you must first configure an IP address and subnet mask for the switch. You can configure the IP information using any of the following:

- DHCP
- Terminal interface through the EIA-232 port

After you establish in-band connectivity, you can change the IP information using any of the following:

- Terminal interface through the EIA-232 port
- Terminal interface through Telnet
- SNMP-based management
- Web-based management

To configure an IPv4 management VLAN or reset the IPv4 management VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Management>Management Interfaces > IPv4 Management VLAN Configuration**.
   The IPv4 Management VLAN Configuration page displays.
3. In the **Management VLAN ID** field, specify the management VLAN ID of the switch.
   The IPv4 management VLAN is used for management of the switch. You can configure any value in the range of 1 to 4093.
4. Select the Routing Mode **Enable** or **Disable** radio button to enable or disable global routing on the switch.
   The default value is Enable.

5. Select the Configuration Method **DHCP** or **Manual** radio button to specify the switch startup action:
   - **DHCP**: The switch requests IP address information from a DHCP server.
   - **Manual**: The switch loads the IP address information that you specify:
     - **IP Address**: Specify the IP address of the interface. The default value is 169.254.100.100.
     - **Subnet Mask**: Specify the IP subnet mask for the interface. This is also referred to as the subnet/network mask and defines the portion of the interface's IP address that is used to identify the attached network. The default value is 255.255.0.0.
     - **Gateway**: Specify the gateway for the interface. The default value is 0.0.0.0.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

NOTE: If you need to reset the IPv4 management VLAN, in the Reset IPv4 Management Interface section of the page, select the **Set Management Interface to Default** check box and click the **Apply** button. Doing so deletes the port-based IPv4 management interface configuration and sets the IPv4 management interface back to the default VLAN 1.

The Current IPv4 Management Interface Status section displays the following view-only fields.

Table 17.IPv4 management interface status information

Field	Description
Management Interface	Displays the current IPv4 management VLAN
Link State	Indicates whether the link status is up or down
Routing Interface Status	Indicates whether the link status is up or down for the management interface
MAC Address	The MAC address assigned to the management interface
IP Address Configuration Method	Indicates whether the IP address configuration method is DHCP or manual
IP Address	The IP address of the management interface
Subnet Mask	The IP subnet mask for the management interface
Gateway	The specified default gateway for the management interface

Configure an IPv4 management interface

You can restrict IPv4 management to one specific interface. By default, you can use any of the interfaces as an IPv4 management interface.

To configure an IPv4 management interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Management>Management Interfaces > IPv4 Management Interface Configuration**.
   The IPv4 Management Interface Configuration page displays.
3. From the **Interface** menu, select the interface that must become the IPv4 management interface.
4. Select the **Set Management Interface** option check box so that you can configure the interface IP address settings
5. Select a Configuration Method radio button:
   - **DHCP**: The switch obtains the management interface IP address settings from a DHCP server in the network.
   - **Manual**: Configure the management interface IP address settings manually:
     - **IP Address**: Specify the IP address for the management interface.
     - **Subnet Mask**: Specify the subnet mask for the management interface.
     - **Gateway**: Specify the IP address of the default gateway for the management interface.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

NOTE: If you need to reset the IPv4 management interface, in the Reset IPv4 Management Interface section of the page, select the Set Management Interface to Default check box and click the **Apply** button. Doing so deletes the IPv4 management interface configuration and sets the IPv4 management interface back to the default setting.

The Current IPv4 Management Interface section displays the following view-only fields.

Table 18.IPv4 management interface status information

Field	Description
Management Interface	Displays the current IPv4 management interface
Link State	Indicates whether the link status is up or down
Routing Interface Status	Indicates whether the link status is up or down for the management interface
MAC Address	The MAC address assigned to the management interface
IP Address Configuration Method	Indicates whether the IP address configuration method is DHCP or manual
IP Address	The IP address of the management interface
Subnet Mask	The IP subnet mask for the management interface
Gateway	The specified default gateway for the management interface

IPv6 management interfaces and VLANs

The main UI provides separate options for IPv6 interface and port-based IP management.
If you configure port-based IP management, VLAN-based IP management is disabled. Similarly, if you configure VLAN-based IP management, port-based IP management is disabled.
By default, the source interface for applications is VLAN 1. If you change the IPv6 management VLAN and port to a non-default value, the source interface to the VLAN 1 default VLAN and port are also automatically changed.
You can access the switch over the main UI by one of the following methods:

- **IPv6 service port**: The service port is a dedicated Ethernet port for out-of-band management of the switch. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.
- **Management VLAN interface**: The management VLAN is the logical interface used for in-band connectivity with the switch over any of the switch's network interfaces. To access the

switch over a network you must first configure an IP address and subnet mask for the switch. To configure an IPv6 management VLAN, you can either use the same VLAN that is used for IPv4 management or a different VLAN. The switch does not provide a default IPv6 management VLAN. If you want to use one, you must create it manually.

- **Optional single management interface**: By default, you can use any of the switch's network interfaces as the management interface. However, you can restrict IPv6 management to one specific network interface.

Configure an IPv6 service port

You can configure IPv6 network information on the service port. The service port is a dedicated Ethernet port for out-of-band management of the switch Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.

To configure the IPv6 service port:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Management>Management Interfaces > IPv6 Service Port Configuration**.
   The IPv6 Service Port Configuration page displays.
3. Select the IPv6 mode **Enable** or **Disable** radio button.
   This specifies the IPv6 administrative mode on the service port.
4. Select the Service Port Configuration Protocol **None** or **DHCP** radio button.
   This specifies whether the device acquires network information from a DHCPv6 server. Selecting **None** disables the DHCPv6 client on the service port.
   If you select the **DHCP** radio button, the DHCPv6 Client DUID field displays the client identifier used by the DHCPv6 client when sending messages to the DHCPv6 server.
5. Select the IPv6 Stateless Address AutoConfig mode **Enable** or **Disable** radio button:
   - **Enable**: The service port can acquire an IPv6 address through IPv6 Neighbor Discovery Protocol (NDP) and through the use of router advertisement messages.
   - **Disable**: The service port does not use the native IPv6 address autoconfiguration feature to acquire an IPv6 address.

   This sets the IPv6 stateless address autoconfiguration mode on the service port.

6. To configure the IPv6 gateway, do the following:
   a. Select the **Change IPv6 Gateway** check box.
      The IPv6 gateway is the default gateway for the IPv6 service port interface.
   b. In the **IPv6 Gateway** field, specify the default gateway for the IPv6 service port interface.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Manage IPv6 addresses for the IPv6 service port

You can add a static IPv6 address that is specific to the IPv6 service port. You can also remove an IPv6 address that you no longer need for the IPv6 service port.

To add or remove an IPv6 address for the IPv6 service port:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Management Interfaces > IPv6 Service Port Configuration**.
   The IPv6 Service Port Configuration page displays.
   The Add/Delete IPv6 Address table lists the static IPv6 addresses that you manually added to the service port interface.
3. To add an IPv6 address, in the Add/Delete IPv6 Address section, do the following:
   a. In the **IPv6 Address** field, specify the IPv6 address that you want to add.
   b. From the **EUI FLAG** menu, select **True** to enable the Extended Universal Identifier (EUI) flag for the IPv6 address, or select **False** to omit the flag.
   c. Click the **Add** button.
      The IPv6 address is added to the IPv6 service port.
4. To remove an IPv6 address, in the Add/Delete IPv6 Address section, do the following:
   a. Select the check box for the IPv6 address that you want to remove.
   b. Click the **Delete** button.
      The IPv6 address is removed from the IPv6 service port.
5. To save the settings to the running configuration, click the **Save** icon.

Configure an IPv6 management VLAN

For you to manage the switch over the main UI, you must define the management IP address. A management VLAN interface is created by default and it gets an IP address if a DHCP server is present. If the management VLAN interface cannot get an IP address, the automatically assigned fallback IPv4 address is 169.254.100.100 and the subnet mask is 255.255.0.0.
A management VLAN is used as the default source interface for the syslog, the message log, an SNMP client, and so on. The network interface is disabled by default.
The management VLAN is the logical interface used for in-band connectivity with the switch through any of the switch's front panel ports. The configuration parameters associated with the switch's management VLAN do not affect the configuration of the front panel ports through which traffic is switched or routed.
To access the switch over a network you must first configure an IP address and subnet mask for the switch. You can configure the IP information using any of the following:

- DHCP
- Terminal interface through the EIA-232 port

After you establish in-band connectivity, you can change the IP information using any of the following:

- Terminal interface through the EIA-232 port
- Terminal interface through Telnet
- SNMP-based management
- Web-based management

To configure an IPv6 management VLAN, you can either use the same VLAN that is used for IPv4 management or a different VLAN. The switch does not provide a default IPv6 management VLAN. If you want to use one, you must create it manually.

To configure an IPv6 management VLAN or reset the IPv6 management VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Management>Management Interfaces > IPv6 Management VLAN Configuration**.
   The IPv6 Management VLAN Configuration page displays.
3. In the **Management VLAN ID** field, specify the management VLAN ID of the switch.
   The IPv6 management VLAN is used for management of the switch. You can configure any value in the range of 1 to 4093.
4. Select the IPv6 Enable Mode **Enable** or **Disable** radio button to specify he administration mode for the IPv6 management VLAN.

5. Select the Address Autoconfigure Mode **Enable** or **Disable** radio button to specify the address autoconfiguration mode.
   If you enable this mode, the IPv6 network settings (IPv6 prefix and prefix length) are automatically configured for the configured management VLAN interface. By default, for VLAN 1, the mode is enabled so that the settings are automatically configured.

   NOTE: The Address Autoconfigure mode option is available only if unicast routing is globally disabled.

6. Select the Address DHCP Mode **Enable** or **Disable** radio button to specify if the IPv6 settings are obtained from a DHCP server.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

NOTE: If you need to reset the IPv6 management interface, in the Reset IPv6 Management Interface section of the page, select the **Set Management Interface to Default** check box and click the **Apply** button. Doing so deletes the IPv6 management interface configuration.

The Current IPv6 Management Interface Status section displays the following view-only fields.

Table 19.Current IPv6 management interface status information

Field	Description
Management Interface	The IPv6 management interface. (By default, no IPv6 management interface is configured).
Link State	Indicates if the link status is up or down.
IPv6 Routing Interface Status/ Operational Mode	Indicates if the link status is up or down for the IPv6 management interface.
MAC Address	The MAC address that is assigned to the IPv6 management interface.
IPv6 Enable Mode	Indicates if IPv6 is enabled or disabled on the IPv6 management interface.
IPv6 Routing Mode	Indicates if IPv6 routing is enabled or disabled on the IPv6 management interface.
DHCPv6 Client Mode	Indicates if the DHCPv6 Client Mode is enabled or disabled on the IPv6 management interface.
IPv6 Default Gateway	The IPv6 address of the default gateway of the switch.
IPv6 Next Hop Interface	The IPv6 address of the next hop interface of the switch.
IPv6 Prefix/Prefix Length	The IPv6 prefix and prefix length that is used for the IPv6 management interface.
EU164	The EUI-64 flag of the IPv6 address on the IPv6 management interface.
Current State	The current state of the IPv6 address on the IPv6 management interface.

Manage IPv6 addresses for the IPv6 management VLAN

You can add a static IPv6 address that is specific to the IPv6 management VLAN. You can also remove an IPv6 address that you no longer need for the IPv6 management VLAN.

To add or remove an IPv6 address for the IPv6 management VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Management Interfaces > IPv6 Management VLAN Configuration**.
   The IPv6 Management VLAN Configuration page displays.
   If the Address Autoconfigure Mode is enabled, the autoconfigured IPv6 prefix and prefix length are displayed in the IPv6 VLAN Interface Configuration section.
3. To add an IPv6 address, in the IPv6 VLAN Interface Configuration section, do the following:
   a. In the **IPv6 Prefix/Prefix Length** field, specify the IPv6 address and prefix length that you want to add.
   b. From the **EUI64** menu, select **True** to enable the Extended Universal Identifier (EUI) flag for the IPv6 address, or select **False** to omit the flag.
   c. Click the **Add** button.
      The IPv6 address is added to the IPv6 management VLAN.
4. To remove an IPv6 address, in the IPv6 VLAN Interface Configuration section, do the following:
   a. Select the check box for the IP address and prefix that you want to remove.
   b. Click the **Delete** button.
      The IPv6 address is removed from the IPv6 management VLAN.
5. To save the settings to the running configuration, click the **Save** icon.

Manage the IPv6 default route addresses for the IPv6 management VLAN

You can add an IPv6 default route address that is specific to the IPv6 management VLAN. You can also remove an IPv6 default route address that you no longer need for the IPv6 management VLAN.

To add or remove an IPv6 default route address for the IPv6 management VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **System > Management > Management Interfaces > IPv6 Management VLAN Configuration**.
   The IPv6 Management VLAN Configuration page displays.
3. To add an IPv6 default route address, in the IPv6 Default Route Configuration section, do the following:
   a. Select the **Change IPv6 Default Route** check box.
   b. In the **IPv6 Default Route Address** field, specify the IPv6 address and prefix length that you want to add.
   c. Click the **Add** button.
      The IPv6 default route address is added to the IPv6 management VLAN.
4. To remove an IPv6 default route address, in the IPv6 Default Route Configuration section, do the following:
   a. Select the check box for the IPv6 default route address that you want to remove.
   b. Click the **Delete** button.
      The IPv6 default route address is removed from the IPv6 management VLAN.
5. To save the settings to the running configuration, click the **Save** icon.

Configure an IPv6 management interface

You can restrict IPv6 management to one specific interface. By default, you can use any of the interfaces as an IPv6 management interface.

To configure an IPv6 management interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Management>Management Interfaces > IPv6 Management Interface Configuration**.
   The IPv6 Management Interface Configuration page displays.
3. From the **Interface** menu, select the interface that must become the IPv6 management interface.
4. Select the **Set Management Interface** option check box so that you can configure the interface settings.
5. Select the Routing Mode **Enable** or **Disable** radio button to enable or disable routing on the IPv6 management interface.
6. Select the IPv6 Enable Mode **Enable** or **Disable** radio button to enable or disable IPv6 on the IPv6 management interface.

7. Select the Stateless Address Autoconfigure Mode **Enable** or **Disable** radio button to enable or disable the stateless address autoconfiguration mode on the IPv6 management interface. You can enable this mode only if unicast routing is globally disabled.
8. Select the DHCPv6 Client Mode **Enable** or **Disable** radio button to enable or disable the DHCPv6 client mode on the IPv6 management interface.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

NOTE: If you need to reset the IPv6 management interface, in the Reset IPv6 Management Interface section of the page, select the **Set Management Interface to Default** check box and click the **Apply** button. Doing so deletes the IPv6 management interface configuration.

The Current IPv6 Management Interface Status section displays the following view-only fields.

Table 20.Current IPv6 management interface status information

Field	Description
Management Interface	The IPv6 management interface (by default, no IPv6 management interface is configured)
Link State	Indicates if the link status is up or down
IPv6 Routing Interface Status/ Operational Mode	Indicates if the link status is up or down for the IPv6 management interface
MAC Address	The MAC address that is assigned to the IPv6 management interface
IPv6 Enable Mode	Indicates if IPv6 is enabled or disabled on the IPv6 management interface
IPv6 Routing Mode	Indicates if IPv6 routing is enabled or disabled on the IPv6 management interface
DHCPv6 Client Mode	Indicates if the DHCPv6 Client Mode is enabled or disabled on the IPv6 management interface
IPv6 Default Gateway	The IPv6 address of the default gateway of the switch
IPv6 Next Hop Interface	The IPv6 address of the next hop interface of the switch
IPv6 Prefix/Prefix Length	The IPv6 prefix and prefix length that is used for the IPv6 management interface
EU164	The EUI-64 flag of the IPv6 address on the IPv6 management interface
Current State	The current state of the IPv6 address on the IPv6 management interface

Manage IPv6 addresses for the IPv6 management interface

You can add a static IPv6 address that is specific to the IPv6 management interface. You can also remove an IPv6 address that you no longer need for the IPv6 management interface.

To add or remove an IPv6 address for the IPv6 management interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Management Interfaces > IPv6 Management Interface Configuration**.
   The IPv6 Management Interface Configuration page displays.
3. To add an IPv6 address, in the IPv6 Management Interface Configuration section, do the following:
   a. Select the **Set Management Interface** check box.
      It is likely that you already did this when you started the configuration of the IPv6 management interface.
   b. In the **IPv6 Prefix** field, specify the IPv6 address that you want to add.
   c. In the **Prefix Length** field, specify the prefix length for the IPv6 address that you want to add.
   d. From the **EUI64** menu, select **Enable** to enable the Extended Universal Identifier (EUI) flag for the IPv6 address, or select **Disable** to omit the flag.
   e. Click the **Add** button.
      The IPv6 address is added to the IPv6 management interface.
4. To remove an IPv6 address, in the IPv6 Management Interface Configuration section, do the following:
   a. Select the check box for the IPv6 that you want to remove.
   b. Click the **Delete** button.
      The IPv6 address is removed from the IPv6 management interface.
5. To save the settings to the running configuration, click the **Save** icon.

Manage the IPv6 default route addresses for the IPv6 management interface

You can add an IPv6 default route address that is specific to the IPv6 management interface. You can also remove an IPv6 default route address that you no longer need for the IPv6 management interface.

To add or remove an IPv6 default route address for the IPv6 management interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Management Interfaces > IPv6 Management Interface Configuration**.
   The IPv6 Management Interface Configuration page displays.
3. To add an IPv6 default route address, in the IPv6 Default Route Configuration section, do the following:
   a. Select the **Change IPv6 Default Route** check box.
   b. In the **IPv6 Default Route Address** field, specify the IPv6 address and prefix length that you want to add.
   c. Click the **Add** button.
      The IPv6 default route address is added to the IPv6 management interface.
4. To remove an IPv6 default route address, in the IPv6 Default Route Configuration section, do the following:
   a. Select the check box for the IPv6 default route address that you want to remove.
   b. Click the **Delete** button.
      The IPv6 default route address is removed from the IPv6 management interface.
5. To save the settings to the running configuration, click the **Save** icon.

Time and NTP settings

The switch supports the Network Time Protocol (NTP), which is a protocol for synchronizing the clocks of networked computer systems, primarily when data transfer is handled through the Internet.

Configure the global NTP settings

To configure the global NTP settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Time > Time Configuration**.
   The Time Configuration page displays.
   The Date and Time fields display the current date and time that are detected by the switch.

3. Select the NTP Authentication Mode **Enable** or **Disable** radio button to enable or disable NTP packet authentication on the switch.
   When enabled, an NTP packet from an NTP server is authenticated. The authentication is based on SHA256. If the packet cannot be authenticated, it is dropped, and the switch cannot synchronize its time with the NTP server that sent the packet.
   By default, this mode is enabled.
4. Select the NTP Broadcast Client Mode **Enable** or **Disable** radio button to enable or disable the switch from broadcasting NTP requests to an NTP server.
   Broadcasting an NTP request can still be useful if there is no public NTP server configured on the switch because it allows the switch to synchronize its time with an NTP server that is accessible in the network, as opposed to an external time server, such a public time server. By default, this mode is disabled.
5. In the **Broadcast Delay** field, specify the estimated round-trip delay in microseconds between the NTP client and the NTP broadcast server.
   The range is from 1 to 999999. The default is 3000.
   The delay is the period that is lost between the outgoing request and the incoming response because of network traffic and routing. The switch adds the delay that you specify to the time that it receives from the NTP server so that the switch can calculate the actual time.
6. From the **Source Interface** menu, select the outbound interface.
   The NTP client uses the source IP address of this interface to broadcast NTP queries. By default, the following options display in the menu:
   - **None**: The primary IP address of the originating (outbound) interface
   - **VLAN 1**: The primary IP address of VLAN 1. This is the default selection.
   - **Service Port**: The management port IP address

   Depending on the configuration of your switch, the following options can display:
   - **Another VLAN ID**: The primary IP address of a VLAN other than VLAN 1
   - **Routing interface**: The primary IP address of a routing interface
   - **Routing VLAN**: The primary IP address of a VLAN routing interface
   - **Routing loopback interface**: The primary IP address of a routing loopback interface
   - **Different**: For some features, *Different* can display. This means that the source interface is configured separately.

7. In the **Time Zone Name** field, specify a time zone.
   In the **Offset Hours** and **Offset Minutes** fields, you can optionally specify the number of hours and number of minutes that the time zone is different from Coordinated Universal Time (UTC). The time zone can affect the display of the current system time.

   NOTE: When you use a public NTP time server, the time data that is received from the server is based on UTC, which is the same as Greenwich Mean Time (GMT). This might not be the time zone in which the switch is located but the time zone in which you want the switch to operate.

8. In the **Offset Hours** field, specify the number of hours that the time zone is different from UTC.
   The range is from –12 to 14. The default is 0.
9. In the **Offset Minutes** field, specify the number of minutes that the time zone is different from UTC.
   The range is from 0 to 59. The default is 0.
10. Click the **Apply** button.
    Your settings are saved.
11. To save the settings to the running configuration, click the **Save** icon.

Display the global NTP status and packet information

You can display the global NTP status and packet information.

To display the global NTP status and packet information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Time > Time Configuration**.
   The Time Configuration page displays.

3. To refresh the page, click the **Refresh** button.
   The NTP Global Status section displays the following view-only fields.

Table 21.NTP global status information

Field	Description
Version	The NTP version that the client supports
Supported mode	The NTP modes that the client supports. Multiple modes can be supported by a client.
Server IP Address	The IP address of the time server for the last received valid packet. If no message was received from any server, an empty string is shown.
Address Type	The address type of the time server address for the last received valid packet
Server Stratum	The claimed stratum of the time server for the last received valid packet
Reference Clock ID	The reference clock identifier of the time server for the last received valid packet
Server mode	The mode of the time server for the last received valid packet
Unicast Server Max Entries	The maximum number of unicast server entries that can be configured for this client
Unicast Server Current Entries	The number of current valid unicast server entries configured for this client

The NTP Information section displays the following view-only fields.

Table 22.NTP information

Field	Description
NTP Software Name	The name of the NTP software application
NTP Software Version	The version of the NTP software application
NTP System Type	The operating system platform information of the NTP system

The NTP Packets section displays the following view-only fields.

Table 23.NTP packets information

Field	Description
NTP In packets	The number of NTP messages that the switch received
NTP Out packets	The number of NTP messages that the switch sent

Field	Description
NTP old version packets	The number of NTP messages for an unsupported NTP version that the switch received
NTP protocol error packets	The number of NTP messages with an NTP protocol error that the switch received

Configure NTP servers

NTP assures accurate time synchronization for the network device clock, up to the millisecond. Time synchronization is performed by a network NTP server. The switch operates as an NTP client only and does not provide time services to other devices.
Time sources are established by stratums. Stratums define the accuracy of the reference clock. The higher the stratum (zero is the highest), the more accurate the clock. The switch receives time from stratum 1 or stratum 0 devices because the switch itself is a stratum 2 device.
The following are examples of stratums:

- **Stratum 0**: A real-time clock is used as the time source, for example, a GPS system.
- **Stratum 1**: A server that is directly linked to a stratum 0 time source is used. Stratum 1 time servers provide primary network time standards.
- **Stratum 2**: The time source is distanced from the stratum 1 server over a network path. For example, a stratum 2 server receives the time over a network link, through NTP, from a stratum 1 server.

Information received from NTP servers is evaluated based on the time level and server type. NTP time definitions are assessed and determined by the following time levels:

- **T1**: Time that the original request was sent by the client.
- **T2**: Time that the original request was received by the server.
- **T3**: Time that the server sent a reply.
- **T4**: Time that the client received the server's reply.

T1 through T4 are used to determine the server time.
The switch can poll unicast server types for the server time. The switch polls for unicast information to detect a server for which the IP address is known. NTP servers that you configure on the switch are the only ones that are polled for synchronization information. This is the preferred method for synchronizing device time because it is the most secure method.
The switch retrieves synchronization information either by actively requesting information or at every poll interval.

Add an NTP server

The switch is preconfigured with public NTP servers but you can also add NTP servers.

To add an NTP server:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Time > NTP Server Configuration**.
   The NTP Server Configuration page displays.
3. From the **Server Type** menu, select the type of NTP address to enter in the address field.
   The address can be an IPv4 address, IPv6 address, or host name (DNS). The default is DNS.
4. In the **Address** field, specify the IP address or the host name of the NTP server.
   This is an IP address or a text string of up to 64 characters, containing the encoded unicast IP address or host name of an NTP server. Unicast NTP requests are sent to this address. If this address is a DNS host name, the host name is resolved into an IP address each time an NTP request is sent to it.
5. In the **Version** field, specify the NTP version that is supported by the switch.
   The range is from 1 to 4. The default is 4.
6. In the **Key ID** field, specify the authentication key number that the switch uses to communicate with the NTP server.
   The range is from 0 to 4294967295.
7. In the **Minimum Poll Interval** field, specify the shortest polling interval in powers of two seconds.
   The range is from 2^4 to 2^{17}. The default is 2^6.
8. In the **Maximum Poll Interval** field, specify the longest polling interval in powers of two seconds.
   The range is from 2^4 to 2^{17}. The default is 2^{10}.
9. From the **Prefer** menu, select **Yes** to assign the NTP server as the preferred NTP server.
   A preferred server reduces the number of times that the switch must change between different NTP servers. The default is No, which means that an NTP server is not a preferred server.
10. From the **Burst** menu, select **Yes** to let the switch send a series of packets instead of a single packet within each synchronization interval.
    If you select No, a single packet is sent, which might cause slower synchronization. The default is Yes.
11. From the **Iburst** menu, select **Yes** to let the switch send a series of packets instead of a single packet within the initial synchronization interval.
    If you select No, a single packet is sent, which might cause slower initial synchronization. The default is Yes.

12. Click the **Add** button.
    The NTP server entry is added.
13. To add additional NTP servers, repeat the previous steps.
14. To save the settings to the running configuration, click the **Save** icon.
    The NTP Server Associations table displays the following information.

Table 24.NTP server associations information

Field	Description
Association Name	The IP address or host name for the NTP association
Reference ID	The reference ID for the NTP association
Server Type	The type of address (IPv4 or IPv6) for the association
Address	The IPv4 or IPv6 address for the association
Offset	The offset time for the association.
Stratum	The stratum version for the association
Jitter	The jitter in milliseconds for the association
Delay	The network delay in milliseconds for the association
Dispersion	The root dispersion (error factor) for the association.
NTP In packets	The number of NTP messages that the switch received for the association
NTP Out packets	The number of NTP messages that the switch sent for the association
NTP protocol error packets	The number of NTP messages with an NTP protocol error that the switch received for the association

Change the settings for an existing NTP server

You can change the setting for an existing preconfigured or custom NTP server.

To change the settings for an existing NTP server:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Time > NTP Server Configuration**.
   The NTP Server Configuration page displays.
3. In the NTP Server Configuration section, select the check box next to the NTP server.

4. Change the settings as needed.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Remove an NTP server

You can remove an NTP server that you no longer need.

To remove an NTP server:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Time > NTP Server Configuration**.
   The NTP Server Configuration page displays.
3. In the NTP Server Configuration section, select the check box next to the NTP server.
4. Click the **Delete** button.
   The server is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure daylight saving time settings

You can configure settings for summer time, which is also known as daylight saving time. Used in some countries around the world, summer time is the practice of temporarily advancing clocks during the summer months. Typically clocks are adjusted forward one or more hours near the start of spring and are adjusted backward in autumn.
To configure the daylight saving time settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Time > Daylight Saving Configuration**.
   The Daylight Saving Configuration page displays.

3. Select one of the following Daylight Saving (DST) radio buttons:
   - **Disable**: Disable daylight saving time.
   - **Recurring**: Daylight saving time occurs at the same time every year. You must manually configure the start and end times and dates for the time shift. Configure the settings that are described in step 4.
   - **Recurring EU**: The system clock uses the standard recurring summer time settings used in countries in the European Union. With this option, the rest of the applicable fields on the page are automatically populated and you cannot change them. Configure the settings that are described in step 4.
   - **Recurring USA**: The system clock uses the standard recurring daylight saving time settings used in the United States. With this option, the rest of the applicable fields on the page are automatically populated and you cannot change them. Configure the settings that are described in step 4.
   - **Non Recurring**: Daylight saving time settings are in effect only between the start date and end date of the specified year. With this option, the summer time settings do not repeat on an annual basis. Configure the settings that are described in step 5.
4. If you select **Recurring**, **Recurring EU**, or **Recurring USA** radio button, configure the fields that are shown in the following table.

Field	Description
Begins At	These fields are used to configure the start values of the day and time. ○ **Week**: Configure the start week. ○ **Day**: Configure the start day. ○ **Month**: Configure the start month. ○ **Hours**: Configure the start hours. ○ **Minutes**: Configure the start minutes.
Ends At	These fields are used to configure the end values of day and time. ○ **Week**: Configure the end week. ○ **Day**: Configure the end day. ○ **Month**: Configure the end month. ○ **Hours**: Configure the end hours. ○ **Minutes**: Configure the end minutes.
Offset	Configure recurring offset in minutes. The valid range is from 1 to 1440 minutes.
Zone	Configure the time zone.

5. If you select the **Non Recurring**, configure the fields that are shown in the following table.

Field	Description
Begins At	These fields are used to configure the start values of the date and time. ○ **Month**: Configure the start month. ○ **Date**: Configure the start date. ○ **Year**: Configure the start year. ○ **Hours**: Configure the start hours. ○ **Minutes**: Configure the start minutes.
Ends At	These fields are used to configure the end values of date and time. ○ **Month**: Configure the end start date. ○ **Date**: Configure the end date. ○ **Year**: Configure the end year. ○ **Hours**: Configure the end hours. ○ **Minutes**: Configure the end minutes.
Offset	Configure the non-recurring offset in minutes. The valid range is from 1 to 1440 minutes.
Zone	Configure the time zone.

6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

View the daylight saving time status

You can display information about the summer time settings and whether the time shift for summer time is currently in effect.

To view the daylight saving time status:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Time > Daylight Saving Configuration**.
   The Daylight Saving Configuration page displays.

3. To refresh the page, click the **Refresh** button.
   The following table displays the view-only daylight saving (DST) status information.

Table 25.Daylight saving status information

Field	Description
Daylight Saving (DST)	The Daylight Saving value, which is one of the following: ○ **Disable** ○ **Recurring** ○ **Recurring EU** ○ **Recurring USA** ○ **Non Recurring**
Begins At	Displays when the daylight saving time begins. This field is not displayed when daylight saving time is disabled.
Ends At	Displays when the daylight saving time ends. This field is not displayed when daylight saving time is disabled.
Offset (in Minutes)	The offset value in minutes.This field is not displayed when daylight saving time is disabled.
Zone	The zone acronym. This field is not displayed when daylight saving time is disabled.
Daylight Saving (DST) in Effect	Displays whether daylight saving time is in effect.

Precision Time Protocol transparent clock

Precision Time Protocol (PTP, IEEE 1588) is a protocol that enables precise synchronization of clocks with a sub-microsecond accuracy across a packet-based network. PTP lets network devices of different precision and resolution synchronize to a grandmaster clock through an exchange of packets across the network. The switch supports a PTP end-to-end transparent clock (TC), which is enabled by default, both globally and at the port level.
Models M4350-8M2V, M4350-16M4V, M4350-24M4X4V, M4350-24F4, M4350-16V4C, M4350-40X4C, M4350-40F4C, and M4350-16C also support the PTP boundary clock and grandmaster clock features. For more information, see Precision Time Protocol boundary clock and grandmaster clock on page 78.

NOTE: The switch itself is not affected by PTP.

NOTE: For M4350 series SMPTE-capable switches, you can configure both PTP-BC and PTP-TC on different ports of the same switch, as long as each PTP type is on a separate VLAN. For a list of M4350 series SMPTE-capable switches, see NETGEAR AV and IT Managed Switches.

Manage the global PTP TC settings

By default, the PTP end-to-end transparent clock (TC) is enabled globally on the switch. You can disable PTP TC globally, in which case the switch does not support PTP pass-through.

To configure the PTP (TC) settings globally:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > PTP TC > Global Configuration**.
   The Global Configuration page displays.
3. Select the Admin Mode **Enable** or **Disable** radio button.
   The default is Disable.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Manage the PTP TC interface settings

On a standalone switch, by default, PTP end-to-end transparent clock (TC) is enabled globally on all interfaces. In a switch stack, by default, PTP TC is disabled for all interfaces. You can select individual interfaces on which you can enable or disable PTP TC. If you disable PTP TC on an interface, the interface does not support PTP pass-through.

To configure the PTP TC settings for one or more interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > PTP TC > Interface Configuration**.
   The Interface Configuration page displays.

3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Configured Mode** menu, select **Enable** or **Disable**.
   The default is Enable.
6. Click the **Apply** button.
   Your settings are saved.
   The Operational Mode field shows whether PTP is enabled or disabled for an interface.
7. To save the settings to the running configuration, click the **Save** icon.

Manage PTP multiple domains

Precision Time Protocol multiple-domain (PTP-MD) allows a single switch to host multiple instances of PTP BC simultaneously. A switch can host an SMPTE profile along with an AES67 profile, with a connected grandmaster for each profile, or acting as a grandmaster for any of the profiles. You can create up to 4 PTP domains using the following profiles:

- **Default**: The default PTP settings are used. This is the default setting.
- **AES67**: A clock profile for connecting audio devices over a common PTP profile.
- **SMPTE-2059-2**: A clock profile for video traffic.
- **AES-R16-2016**: A clock profile that interoperates between SMPTE and AES67 clock profiles.

Multiple PTP Boundary Clock (PTP BC) instances in a switch are distinguished by their domain identifiers.
You must configure the profile with a unique domain ID, or associate it with a unique VLAN.

NOTE: You can associate each PTP BC instance, or domain, with a VLAN. Multiple domains can share a VLAN, or each use a different one, but a domain maps to only one VLAN. The egress port's VLAN membership determines whether PTP control packets leave tagged or untagged, not the domain-to-VLAN association. When each BC instance uses a distinct VLAN carried as a tagged member on a shared trunk port, a physical port or a LAG, that port can carry PTP control packets for multiple BC instances, separated by VLAN.

Precision Time Protocol boundary clock and grandmaster clock

NOTE: The PTP boundary clock and grandmaster clock features are supported on the following models:

- M4350-8M2V
- M4350-16M4V
- M4350-24M4X4V
- M4350-24F4
- M4350-16V4C
- M4350-40X4C
- M4350-40F4C
- M4350-16C

Precision Time Protocol (PTP, IEEE 1588) is a protocol that enables precise synchronization of clocks with a sub-microsecond accuracy across a packet-based network. PTP lets network devices of different precision and resolution synchronize to a grandmaster clock through an exchange of packets across the network.

The switch supports a boundary clock (BC) and grandmaster (GM) clock. By default, the PTP BC mode is disabled. Both of these models are capable of functioning as a PTP GM clock.

The clocks can be defined as follows:

- **Boundary clock (BC)**: A switch that has one port that functions as a time-transmitter (master) in a PTP network and one or more other ports that function as time-receivers. Depending on the local clock priority value and other settings, the switch clock function in a PTP network might change automatically from boundary clock to grandmaster clock.
- **Grandmaster (GM) clock**: A switch that has one port that function as the time source and time-transmitter (master) in a PTP network. The grandmaster clock provides precise time reference for one or more devices that are directly attached to the switch and that synchronize their clocks with the grandmaster clock. Depending on the local clock priority value and other settings, the switch clock function in a PTP network might change automatically from grandmaster clock to boundary clock.

The best master clock algorithm (BMCA) determines the best time-transmitter (master) for a

PTP domain and determines the time-transmitter–to–time-receiver hierarchy with the best time-transmitter as the root.
Note the following about the PTP BC/GM clock feature:

- If a network profile that uses AVB is enabled (see 802.1AS timing and synchronization on page 323), the PTP BC clock (or transparent clock) is automatically disabled and you cannot manually enable the feature. AVB and PTP cannot coexist.
- PTP is not supported on a stack.

NOTE: For M4350 series SMPTE-capable switches, you can configure both PTP BC and PTP TC on different ports of the same switch, as long as each PTP type is on a separate VLAN. For a list of M4350 series SMPTE-capable switches, see NETGEAR AV and IT Managed Switches.

Manage the PTP global configuration settings

You can set a global PTP source address on your switch using the PTP Global Configuration page. The switch uses the configured global PTP source address for PTP messages sent by all PTP domains.
The PTP Domains section lists the PTP domains that you added. For more information about configuring settings for a specific PTP domain, see Manage the PTP domain, boundary clock, and grandmaster clock settings on page 80.

To manage the PTP global configuration settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > PTP > PTP Global Configuration**.
   The PTP Global Configuration page displays.
   In the PTP Global Configuration section, set the IPv4 and IPv6 address that you want the switch to use as the PTP source address for all PTP domains.
3. In the **Source IPv4 Address** field, enter an IPv4 address.
   The switch supports only one global PTP source address. If you do not configure a global PTP source address, the switch uses the management VLAN IP address instead.
4. In the **Source IPv6 Address** field, enter an IPv6 address.
   The switch uses the same priority order for the IPv6 address as for the IPv4 address.
5. In the PTP Domains section, enter the number of the domain that you want to add in the **Domains** field.
   You can enter a domain number from 0 to 127.

6. Click the **Add** button.
   The domain is added to the **PTP Domains** table.
7. To remove a domain, select the check box next to the domain in the **PTP Domains** table, then click the **Delete** button.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Manage the PTP domain, boundary clock, and grandmaster clock settings

By default, the PTP boundary clock and grandmaster clock mode is disabled on the switch. You can enable the mode and configure the settings on your PTP domain.

To manage the PTP domain configuration settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > PTP > PTP Domain Configuration**.
   The PTP Domain Configuration page displays.
3. Select the domain from the **Domain** menu.
4. In the **Local Clock Priority1** field, enter a priority value from 0 to 255. The priority 1 determines which switch that is capable of functioning as a grandmaster clock is elected to be the active grandmaster clock in the network.
5. In the **Local Clock Priority2** field, enter a value from 0 to 255. The priority 2 also determines which switch that is capable of functioning as a grandmaster clock is elected to be the active grandmaster clock in the network.
6. Depending on your network needs, select one of the following radio buttons for a PTP clock profile:
   - **None**: The PTP profile is removed.
   - **Default**: The default PTPv2 settings are used. This is the default setting. **AES67**: A clock profile for connecting audio devices over a common PTP profile.
   - **SMPTE-2059-2**: A clock profile for video traffic.
   - **AES-R16-2016**: A clock profile that interoperates between SMPTE and AES67 clock profiles.

7. In the **Announce Interval** field, enter the period between successive Announce messages in logarithm to base 2 format.
   Announce messages help to elect the best grandmaster clock in the network. The following are the supported ranges and default settings for the PTP clock profiles:
   - **Default**: 0 to 4. The default setting is 1.
   - **AES67**: 0 to 4. The default setting is 1.
   - **SMPTE-2059-2**: –3 to 1. The default setting is 0.
   - **AES-R16-2016**: 0 or 1. The default setting is 0.
8. In the **Announce Timeout**, enter number of Announce messages that are not acknowledged before the switch determines that the grandmaster clock is not transmitting.
   You can enter a number from 2 to 10. The default setting for all PTP clock profiles is 3. For the AES-R16-2016 profile, the only possible value is 3.
9. In the **Sync Interval** field, enter the period between successive Sync messages, in logarithm to base 2 format.
   The following are the supported ranges and default settings for the PTP clock profiles:
   - **Default**: –1 to 1. The default setting is 0.
   - **AES67**: –4 to 1. The default setting is –3.
   - **SMPTE-2059-2**: –7 to –1. The default setting is –3.
   - **AES-R16-2016**: –4 to 1. The default setting is –3.
10. In the **Delay-request Interval** field, enter the period between successive Delay-Request messages from the boundary clock to the grandmaster clock, in logarithm to base 2 format.
    Delay-request messages help to calculate the path delay. The following are the supported ranges and default settings for the PTP clock profiles:
    - **Default**: 0 to 5. The default setting is 0.
    - **AES67**: –3 to 5. The default setting is 0.
    - **SMPTE-2059-2**: –7 to 4. The default setting is –3.
    - **AES-R16-2016**: –4 to 5. The default setting is –3.
11. Select the **IPv4**, **IPv6**, or **L2 Transport** radio button to set the transport protocol for use on the communication path.
12. In the **Domain** field, enter a value from 0 to 127.
    A PTP domain consists of one of more ordinary or boundary clocks that communicate with each other, and a single grandmaster clock. Although more than one switch can be capable of functioning as a grandmaster in the network, only one switch can the active grandmaster clock.

13. From the **PTP Logging Level** menu, select the logging level for the PTP state transitions and status.
    The options are **Emergency**, **Alert**, **Critical**, **Error**, **Warning**, **Notice** (the default setting), **Info**, and **Debug**. Log messages with the selected severity level and all log messages of greater severity are logged. For example, if you select Warning, the logged messages include Warning, Error, Critical, Alert, and Emergency.
14. Select the PTP Path-trace **Enable** or **Disable** radio button.
    A path trace is the route that a PTP announce message takes through the network path of boundary clocks and is tracked through the path trace TLV in the announce message. The path trace sequence contains the clock ID of each boundary clock that an announce message traverses. The default is Disable.
15. Select a Clock Operation Mode radio button:
    - **One-step**: The clock timestamps the Sync message as it exits the port.
    - **Two-step**: The clock timestamps the Sync message as it exits the port and then immediately follows up with another message that captures the egress timestamp of the Sync message. This is the default setting.

    NOTE: Do not mix one-step and two-step clock operating modes in different PTP domains.
16. In the **IPv6 Multicast Destination Scope** field, type the IPv6 multicast destination scope for the boundary or grandmaster clock.
    By default, the IPv6 multicast destination scope is 0xE.
17. Select the **2.0** or **2.1** radio button to set the PTP version.
    The default PTP version is 2.1.
18. In the **PTP VLAN** field, enter a PTP VLAN for a given domain identifier.
19. Click the **Apply** button.
    Your settings are saved.
20. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 26.PTP global configuration information

Field	Description
PTP process running	Indicates if the PTP process is enabled on the switch.
GrandMaster Capable	Indicates if the switch is capable of becoming a grandmaster clock.
Communication Model	This field is always Multicast.
Local Clock Identity	The ID of the clock, using an EUI-64 assigned number.

Manage the PTP domain, boundary clock, and grandmaster interface settings

On a standalone switch, by default, the PTP boundary clock capability is disabled on all interfaces.

To configure the boundary clock capability for one or more interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > PTP > PTP Interface Configuration**.
   The PTP Interface Configuration page displays.
3. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
4. In the **Domain** field, enter a domain.
   You can add up to four PTP domains to an interface. To assign more than one, enter the domain IDs in the **Domain** field, separated by commas.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only field on the page.

Table 27.PTP boundary clock and grandmaster clock interface configuration information

Field	Description
Operational Mode	Indicates if the PTP boundary clock and grandmaster clock capability is enabled or disabled for an interface.

View or clear the PTP boundary clock statistics

You can view or clear the PTP boundary clock statistics for the interfaces.

To view or clear the PTP boundary clock statistics for the interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > PTP > PTP Statistics**.
   The PTP Statistics page displays.
3. In the PTP Domains section, on the **Domain** menu, select a PTP domain to view the detailed information about the status of the PTP boundary clock for that PTP domain.
   The default domain is **127**.
4. Select one or more interfaces by taking one of the following actions:
   - To view or clear the statistics for a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To clear the statistics for multiple interfaces, select the check box associated with each interface.
   - To clear the statistics for all interfaces, select the check box in the heading row.
5. To clear statistics, click the **Clear** button.
   The statistics for the selected interface or interfaces are cleared.

The following table describes the view-only fields on the page.

Table 28.PTP boundary clock and grandmaster clock statistics information

Field	Description
Interface	The interface for which information is displayed.
Announce	
Transmitted	The number of Announce messages transmitted.
Received	The number of Announce messages received.
Timeouts	The number of Announce messages that timed out.
Sync	
Transmitted	The number of Sync messages transmitted.
Received	The number of Sync messages received.
Mismatched	The number of Sync messages that were mismatched.
Timeouts	The number of Sync messages that timed out.
Followup	
Transmitted	The number of Followup messages transmitted.
Received	The number of Followup messages received.
Mismatched	The number of Followup messages that were mismatched.
Delay Req	
Transmitted	The number of Delay Req messages transmitted.
Received	The number of Delay Req messages received.
Delay Res	
Transmitted	The number of Delay Res messages transmitted.
Received	The number of Delay Res messages received.
Signalling	
Transmitted	The number of Signalling messages transmitted.
Received	The number of Signalling messages received.
Management	
Transmitted	The number of Management messages transmitted.

Field	Description
Received	The number of Management messages received.

View detailed information about the PTP boundary clock

You can view detailed information about the status of the PTP boundary clock and PTP domains on the switch.

To view detailed information about the status of the PTP boundary clock and PTP domains:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > PTP > PTP Details**.
   The PTP Details page displays.
3. In the PTP Domains section, on the **Domain** menu, select a PTP domain to view the detailed information about the status of the PTP boundary clock for that PTP domain.
   The default domain is **127**.

The following table describes the view-only fields on the page.

Table 29.PTP boundary clock and grandmaster clock detailed information

Field	Description
PTP Parent, Parent Clock	
Clock Identity	The clock identity value of the parent clock.
Offset	The offset scaled log variance of the parent clock.
Clock Phase Change Rate	The phase change rate value of the parent clock.
PTP Parent, Grandmaster Clock	
Clock Identity	The clock identity value of the grandmaster clock.
Clock Class	The clock class value of the grandmaster clock.
Clock Accuracy	The clock accuracy value of the grandmaster clock.
Offset	The offset scaled log variance of the grandmaster clock.
Priority1	The configured priority 1 value of the grandmaster clock.
Priority2	The configured priority 2 value of the grandmaster clock.
PTP Boundary-clock	
PTP Device Type	The PTP device type, which can be ordinary clock, transparent clock, or boundary clock.
PTP Device Profile	The device profile, which is the profile for a specific device in a specific network.
PTP Clock Role Type	The boundary clock role can be master clock or time-receiver clock.
Clock Operation Mode	The configured PTP boundary clock operation mode, which can be One-step or Two-step.
Parent Clock Operation Mode	The parent clock operation mode, which can be N/A for master clock or One-step or Two-step for time-receiver clock. If the status is not known, the filed displays Not available.
Clock Domain	The configured domain to which the boundary clock belongs.
Number Of PTP Ports	The number of PTP ports that are available for PTP operation.
Clock Class	The clock class value of the boundary clock.
Clock Accuracy	The clock accuracy value of the boundary clock.
Offset	The offset scaled log variance of the boundary clock.

Field	Description
Steps Removed	The number of hops that the grandmaster clock is removed from the PTP boundary clock.
Mean Path Delay	The calculated end-to-end delay, which is not applicable to PTP. The field displays 0.0.
PTP Time-properties	
Current UTC Offset	The Coordinated Universal Time (UTC) offset value.
Leap59	1 indicates that the last minute of the day contains 59 seconds. 0 indicates that the time scale is not PTP.
Leap61	1 indicates that the last minute of the day contains 61 seconds. 0 indicates that time scale is not PTP.
Current UTC Offset Valid	1 indicates that the Current UTC offset value, Leap59 value, and Leap61 value are known to be correct. 0 indicates that these values are not known or not known to be correct.
PTP Timescale	1 indicates that the PTP time scale is traceable to a primary reference. 0 indicates that the PTP time scale is not traceable.
Time Traceable	1 indicates that the time scale is traceable to a primary reference. 0 indicates that the time scale is not traceable.
Frequency Traceable	1 indicates that the frequency determining the time scale is traceable to a primary reference. 0 indicates that the frequency determining the time scale not traceable.
Time Source	The source of the time that is used by the GM PTP instance.
PTP Path-trace	
Clock Identity	A path trace is the route that a PTP announce message takes through the network path of boundary clocks and is tracked through the path trace TLV in the announce message. The path trace sequence contains the clock identity of each boundary clock that an announce message traverses.

Domain Name System

You can configure information about Domain Name System (DNS) servers that the network uses and how the switch operates as a DNS client.

Configure the global DNS settings and add a DNS server

You can configure the global DNS settings and DNS server information.

To configure the global DNS settings and add a DNS server:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > DNS > DNS Configuration**.
   The DNS Configuration page displays.
3. Select the DNS Status **Disable** or **Enable** radio button:
   - **Enable**: Allows the switch to send DNS queries to a DNS server to resolve a DNS domain name. The default value is Enable.
   - **Disable**: Prevents the switch from sending DNS queries.
4. In the **DNS Default Name** field, enter the name that must be includes in DNS queries.
   When the switch looks up on an unqualified host name, this field provides the domain name. For example, if the default domain name is netgear.com and you enter *test*, then test is changed to test.netgear.com to resolve the name). The maximum length of the name is 255 characters.
5. In the **Retry Number** field, specify the number of times that the switch must retry to send DNS queries to the DNS server.
   Enter a number from 0 to 100. The default is 2.
6. In the **Response Timeout (secs)** field, specify the time in seconds that the switch must wait for a response to a DNS query.
   Enter a number from 0 to 3600. The default is 3.
7. From the **Source Interface** menu, select the interface that must be the source interface for DNS, which means that the primary IP address of the interface is used as the source address.
   By default, the following options display in the menu:
   - **None**: The primary IP address of the originating (outbound) interface
   - **VLAN 1**: The primary IP address of VLAN 1. This is the default selection.
   - **Service Port**: The management port IP address

   Depending on the configuration of your switch, the following options can display:
   - **Another VLAN ID**: The primary IP address of a VLAN other than VLAN 1
   - **Routing interface**: The primary IP address of a routing interface
   - **Routing VLAN**: The primary IP address of a VLAN routing interface
   - **Routing loopback interface**: The primary IP address of a routing loopback interface
   - **Different**: For some features, *Different* can display. This means that the source interface is configured separately.

8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.
10. To add a DNS server to which the switch sends DNS queries, do the following:
    a. In the **DNS Server Address** field in the DNS Server Configuration table, enter an IP address in standard IPv4 or IPv6 dot notation.
    b. Click the **Add** button.
       The server is added to the table. You can specify up to eight DNS servers. The precedence is set in the order that you add the servers.
11. To save the settings to the running configuration, click the **Save** icon.

The following table displays the view-only DNS server information.

Table 30.DNS server configuration information

Field	Description
Serial No	The sequence number of the DNS server.
Preference	The preference of the DNS server. The preference is determined by the order in which you add the servers.

Remove a DNS server

You can remove a DNS server that you no longer need.

To remove a DNS server:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > DNS > DNS Configuration**.
   The DNS Configuration page displays.
3. In the DNS Server Configuration table, select the check box for the DNS server.

   CAUTION: If you do not select a DNS server, all DNS servers are removed after you click the **Delete** button.

4. Click the **Delete** button.
   The DNS server is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure and view host name-to-IP address information

You can manually map host names to IP addresses or view dynamic host mappings.

Add a static entry to the dynamic host mapping table

To add a static entry to the local dynamic host mapping table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > DNS > Host Configuration**.
   The Host Configuration page displays.
3. In the **Host Name** field, specify the static host name.
   The maximum length of the name is 255 characters.
4. In the **IP Address** field, enter the IP address to associate with the host name.
5. Click the **Add** button.
   The entry displays in the Dynamic Host Mapping table.
6. To save the settings to the running configuration, click the **Save** icon.

Remove an entry from the dynamic host mapping table

If you no longer an entry, you can remove it from the dynamic host mapping table.

To remove an entry from the dynamic host mapping table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > DNS > Host Configuration**.
   The Host Configuration page displays.
3. Select the check box next to the entry that you want to remove.
4. Click the **Delete** button.
   The entry is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Maintain the dynamic host mapping table and view the entries

You can change the host name or IP address in an entry of the dynamic host mapping table, view all entries, or clear all entries.

To change the host name or IP address in an entry of the dynamic host mapping table, view all entries, or clear all entries:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > DNS > Host Configuration**.
   The Host Configuration page displays.
3. Select the check box next to the entry to update.
4. Enter the new information in the appropriate field.
5. Click the **Apply** button.
   Your settings are saved.
6. To clear all the dynamic host name entries from the list, click the **Clear** button.
7. To save the settings to the running configuration, click the **Save** icon.
   The dynamic host mapping table shows host name-to-IP address entries that the switch learned. The following table describes the dynamic host fields.

Table 31.Dynamic Host Mapping information

Field	Description
Host	Lists the host name that you assign to the specified IP address.
Total	Time since the dynamic entry was first added to the table.
Elapsed	Time since the dynamic entry was last updated.
Type	The type of the dynamic entry.
Addresses	Lists the IP address associated with the host name.

Switch database management template

A switch database management (SDM) template determines the maximum resources the switch can use for various features. By selecting an SDM template based on how you use the switch in your network, you set the switch to use a specific scaling factor, which determines the allocation of resources. That is, by selecting an SDM template, you reallocate switch resources to support a specific combination of features based on your network requirements.
The main UI provides more information about the templates (see the steps in this task).
The default SDM template is IPv4-Basic. You can change the SDM template for the switch.

To change the SDM template

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > SDM Template Preference**.
   The SDM Template Preference page displays.
3. From the **SDM Next Template ID** menu, select a template:
   - **IPv4-Basic** (which is the same as **Default**): Template for an IPv4 environment with support for basic IPv6. This is the default template.
   - **IPv6-Basic**: Template for an IPv6 environment with support for basic IPv4.
   - **IPv4-BasicPlus**: Template for an IPv4 environment with emphasis on IPv4 unicast and multicast routing and support for IPv6. This template can be used for any M4350 switch model except model M4350-24G4XF.
   - **IPv6-BasicPlus**: Template for an IPv6 environment with emphasis on IPv6 unicast and multicast routing and support for IPv4. This template can be used for any M4350 switch model except model M4350-24G4XF.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

6. To activate the new template, restart the switch.
   The following table displays Summary information.

Table 32.SDM template summary information

Field	Description
SDM Template	Identifies the template
ARP Entries	The maximum number of entries in the Address Resolution Protocol (ARP) cache for routing interfaces
IPv4 Unicast Routes	The maximum number of IPv4 unicast forwarding table entries
IPv6 NDP Entries	The maximum number of IPv6 Neighbor Discovery Protocol (NDP) cache entries
IPv6 Unicast Routes	The maximum number of IPv6 unicast forwarding table entries
ECMP Next Hops	The maximum number of next hops that can be stored in the IPv4 and IPv6 unicast forwarding tables
IPv4 Multicast Routes	The maximum number of IPv4 multicast forwarding table entries
IPv6 Multicast Routes	The maximum number of IPv6 multicast forwarding table entries
Maximum VLAN Entries	The maximum number of VLAN entries, which is 4093 for any SDM template

Green Ethernet settings

You can configure the green Ethernet features to reduce power consumption globally, or on one or more interfaces, by transitioning to low power mode when links are idle.

Configure the global green Ethernet settings

You can configure the global green Ethernet settings.

To configure the global green Ethernet settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Green Ethernet > Green Ethernet Configuration**.
   The Green Ethernet Configuration page displays.

3. Select the Auto Power Down Mode **Disable** or **Enable** radio button.
   By default, this mode is disabled. When a port link is down, the underlying physical layer goes down for a short period and then checks for port link pulses again so that auto-negotiation remains possible. In this way, the switch saves power when no link partner is present for the port.
4. Select the EEE Mode **Disable** or **Enable** radio button.
   By default, this mode is disabled. Energy Efficient Ethernet (EEE) combines the MAC with a family of physical layers that support operation in a low power mode. It is defined by the IEEE 802.3az standard. Lower power mode enables both the send and receive sides of the link to disable some functionality for power savings when lightly loaded. Transition to low power mode does not change the link status. Frames in transit are not dropped or corrupted in transition to and from a port in low power mode. Transition time is transparent to upper layer protocols and applications.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Configure green Ethernet interface settings

You can configure green Ethernet settings for individual interfaces.

To configure the green Ethernet interface settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Green Ethernet > Green Ethernet Interface Configuration**.
   The Green Ethernet Interface Configuration page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.

4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Auto Power Down mode** menu, select **Enable** or **Disable**.
   By default, this mode is disabled for the interface. When a port link is down, the underlying physical layer goes down for a short period and then checks for port link pulses again so that auto-negotiation remains possible. In this way, the switch saves power when no link partner is present for the port. If the Auto Power Down mode is not supported for the interface, N/A is displayed.
6. From the **EEE mode** menu, select **Enable** or **Disable**.
   By default, this mode is disabled for an interface. Energy Efficient Ethernet (EEE) combines the MAC with a family of physical layers that support operation in a low power mode. It is defined by the IEEE 802.3az standard. Lower power mode enables both the send and receive sides of the link to disable some functionality for power savings when lightly loaded. Transition to low power mode does not change the link status. Frames in transit are not dropped or corrupted in transition to and from a port in low power mode. Transition time is transparent to upper layer protocols and applications.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Configure and display detailed green Ethernet settings for an interface

You can select a single interface and configure and display detailed green Ethernet settings.

To configure and display detailed green Ethernet settings for a single interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Green Ethernet > Green Ethernet Details**.
   The Green Ethernet Details page displays.
3. From the **Interface** menu, select the interface.

4. From the **Energy Detect Admin Mode** menu, select to enable or disable this mode.
   By default, this mode is disabled for the interface. When a port link is idle, the underlying physical layer goes into low power idle for a short period of time, and then checks for port link pulses again so that auto-negotiation remains possible. In this way, the switch saves power when no link partner is present for the port.
5. From the **EEE Admin Mode** menu, select to enable or disable this mode.
   With Energy Efficient Ethernet (EEE) mode enabled, the port transitions to low power mode during a link idle condition. The default value is Disabled.
6. In the **EEE Transmit Idle Time** field, enter the period that the interface must remain idle (without receiving packets for transmission) before entering the low-power idle (LPI) state. The range is from 600 to 4294967295 seconds. The default is 600 seconds.
7. In the **EEE Transmit Wake Time** field, enter the period that the interface remains in the LPI state before returning to the active state after it receives a packet for transmission.
   The range is from 8 to 65535 seconds. The default is 17 seconds.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 33.Green Ethernet local device information

Field	Description
Cumulative Energy Saved on this port due to Green mode(s) (Watts * Hours)	The cumulative energy saved on the port because of the green modes that are enabled on the port in watts * hours.
Energy Detect Operational Status	The operational state (Active or Inactive) of the Energy Detect mode.
Reason	The reason for the operational state that is displayed in the previous field.
Tw_sys_tx (uSec)	The value of 'Tw_sys' that the port can support.
Tw_sys_tx Echo (uSec)	The link partner's 'transmit Tw_sys' that the port uses to compute the 'Tw_sys' that it requests from its link partner.
Tw_sys_rx (uSec)	The value of 'Tw_sys' that the port requested from its link partner.
Tw_sys_rx Echo (uSec)	The link partner's 'receive Tw_sys' that the port uses to compute the 'Tw_sys' that it can support.
Fallback Tw_sys (uSec)	The value of the 'fallback Tw_sys' that the port requested from its link partner.
Tx_dll_enabled	The initialization status of the transmit EEE transmit Data Link Layer (DLL) function on the port.
Tx_dll_ready	The transmit DLL status, which indicates if the 'tx system' initialization is complete and if the port is ready to receive or update LLDPDUs that contain EEE TLVs.
Rx_dll_enabled	The status of the EEE capability negotiation on the port.
Rx_dll_ready	The receive DLL status, which indicates if the 'rx system' initialization is complete and if the port is ready to receive or update LLDPDUs that contain EEE TLVs.
Time Since Counters Last Cleared	The time since the counters were reset.

Display green Ethernet information for the link partner of an interface

The switch can detect green Ethernet information for the link partner of an interface, that is for the remote device.

To display the green Ethernet information for the link partner of an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Green Ethernet > Green Ethernet Details**.
   The Green Ethernet Details page displays.
3. Scroll down to the Remote Device Information section.
4. From the **Interface** menu, select the interface.
   The following table describes the view-only fields on the page.

Table 34.Green Ethernet remote device information

Field	Description
Remote ID	The identifier that is assigned to the link partner of the port.
Remote Tw_sys_tx (uSec)	The value of 'Tw_sys' that the link partner can support.
Remote Tw_sys_tx Echo (uSec)	The value of 'transmit Tw_sys' that the link partner returns to the port.
Remote Tw_sys_rx (uSec)	The value of 'Tw_sys' that the link partner requests from the port.
Remote Tw_sys_rx Echo (uSec)	The value of 'receive Tw_sys' that the link partner returns to the port.
Remote Fallback Tw_sys (uSec)	The value of 'fallback Tw_sys' that the link partner returns to the port.

Display the green Ethernet statistics summary

The green Ethernet statistics summary displays information about both the switch and the interfaces.

To display the green Ethernet statistics summary:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **System > Management > Green Ethernet > Green Ethernet Summary**.
   The Green Ethernet Summary page displays.

3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields in the Green Ethernet Statistics Summary section.

Table 35.Green Ethernet statistics summary information

Field	Description
Current Power Consumption (mW)	The power consumption by all ports on the switch in mWatts (mW).
Percentage Power Saving (%)	The percentage of power saved on all ports on the switch when green Ethernet mode is enabled.
Cumulative Energy Saving (W*H)	The cumulative energy saved on the switch in (watts * hour) when all green features are enabled.

The following table describes the view-only fields in the Green Ethernet Feature Summary section.

Table 36.Green Ethernet feature summary information

Field	Description
Unit	1, or the unit ID for a stacked switch
Green Features supported on this unit	The green features that are supported on the switch: ○ Energy-Detect ○ EEE (Energy Efficient Ethernet) ○ LPI-History (EEE low power idle history) ○ LLDP-Cap-Exchg (EEE LLDP capability exchange) ○ Pwr-Usg-Est (Power usage estimates).

The following table describes the view-only fields in the Green Ethernet Interface Summary section.

Table 37.Green Ethernet interface summary information

Field	Description
Interface	The interface for which information is displayed.
Energy Detect Admin Mode	Indicates the status of the mode (Enable or Disable). If the mode is enabled and a port link is down, the underlying physical layer goes down for a short period and then checks for port link pulses again so that auto-negotiation remains possible. In this way, the switch saves power when no link partner is present for the port.
Energy Detect Operational	The operational state (Active or Inactive) of the Energy Detect

Field	Description
Status	mode.
EEE Admin mode	Indicates the status of the mode (Enable or Disable). If this mode is enabled, the port transitions to low power mode during a link idle condition.

Bonjour settings

A Mac that supports Bonjour can discover the switch in the network so that you can find the switch IP address and log in to the main UI of the switch. Bonjour is enabled by default. You can disable Bonjour for security reasons.

Enable or disable Bonjour

To enable or diable Bonjour:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Bonjour >Bonjour Configuration**.
   The Bonjour Configuration page displays.
3. Select one of the following radio buttons:
   - **Enable**. Bonjour is enabled. This is the default setting.
   - **Disable**. Bonjour is disabled.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Display Bonjour information

To display Bonjour information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > Bonjour > Bonjour Details**.
   The Bonjour Details page displays.
   The Bonjour Administration Mode field displays whether Bonjour is enabled or disabled.

3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 38.Bonjour Published Services

Field	Description
Service Name	The Bonjour service names in the switch.
Type	The Bonjour service type names in the switch.
Domain	The Bonjour service domain in the switch.
Port	The Bonjour service port number.
TXT Data	The Bonjour service text.

Enable or disable the TFTP server capability

As a security measure, the Trivial File Transfer Protocol (TFTP) server capability is disabled by default on the switch. You can enable the TFTP server capability so that the switch can reach a TFTP server if one is configured.

To enable or disable the TFTP server capability:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > TFTP Server**.
   The TFTP Server Configuration page displays.
3. Select the Admin Mode **Enable** or **Disable** radio button.
   The default is Disable.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

IPv4 DHCP server

You can configure settings for an IPv4 DHCP server, DHCP pools, and DHCP bindings. You can also view DHCP statistics and, if they occur, DHCP conflicts.

Configure a DHCP server

To configure a DHCP server:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPServer >DHCP Server Configuration**.
   The DHCP Server Configuration page displays.
3. Select the Admin Mode **Disable** or **Enable** radio button.
   This setting specifies whether the DHCP service is enabled or disabled.
   The default is Disable.
4. In the **Ping Packet Count** field, specify the number of packets a server sends to a pool address to check for duplication as part of a ping operation.
   The value can be 0 or from 2 to 10. Setting the value to 0 disables the function. The default is 2.
5. Select the Conflict Logging Mode **Disable** or **Enable** radio button.
   This setting specifies whether conflict logging on a DHCP server is enabled or disabled. The default is Enable.
6. Select the BootP Automatic Mode **Disable** or **Enable** radio button.
   This setting specifies whether BootP for dynamic pools is enabled or disabled. The default value is Disable.
7. To exclude addresses, do the following:
   a. In the **IP Range From** field, enter the lowest address in the range or a single address to be excluded.
   b. In the **IP Range To** field, to exclude a range, enter the highest address in the range. To exclude a single address, enter the same IP address as specified in the **IP Range From** field, or leave it as 0.0.0.0.
   c. Click the **Add** button.
      The addresses that must be excluded are added to the switch.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Add a DHCP class for use with DHCP Option 82

If you let the switch function as a Layer 3 DHCP relay, you can add DHCP classes and configure them for use with DHCP Option 82 information. This is a security feature that lets the DHCP relay detect the DHCP Option 82 information in a DHCP request and forward requests only from trusted devices. For more information about using the switch as a Layer 3 DHCP relay, see Configure the global DHCP relay settings and display the relay statistics on page 115.

To add a DHCP class for use with DHCP Option 82:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPServer >DHCP Class Create**.
   The DHCP Server Class Create page displays.
3. In the **Class Name** field, enter an alphanumerical name of up to 31 characters.
4. Click the **Add** button.
   Your settings are saved. The class is added. You can now configure the class.
5. To save the settings to the running configuration, click the **Save** icon.

Configure a DHCP class

For each DHCP class that you add, you can configure the remote ID and circuit ID suboptions. These are suboptions of DHCP Option 82 and limit the devices that can receive an IP address from the DHCP server.

To configure a DHCP class:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPServer >DHCP Class Configuration**.
   The DHCP Class Configuration page displays.
3. From the **Class Name** menu, select the class.
4. In the **Remote ID** field, specify the ID of the remote host from which the DHCP request can be accepted.
   The remote ID can be a user name, caller ID associated with a telephone number, caller ATM address, modem ID, IP address, and so on.

5. In the **Circuit ID** field, specify the identifier of the circuit from which the DHCP requests can be accepted.
   The circuit ID can also be an interface or VLAN number.
6. Click the **Add** button.
   Your settings are saved. The class is added.
7. To add another set of suboptions for the same class, repeat the previous three steps.
8. To save the settings to the running configuration, click the **Save** icon.

Manage DHCP pools

You can set up and manage different types of pools of IP addresses that the DHCP server can assign.

Create a DHCP pool

You can create a DHCP pool and configure multiple settings for the pool.

To create a DHCP pool:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Services > DHCP Server > DHCP Pool Configuration**.
   The DHCP Pool Configuration page displays.
3. From the **Pool Name** menu, select **Create**.
   The **Pool Name** field displays.
4. In the **Pool Name** field, type a name for the new pool.
   The name can be up to 31 characters in length.

5. From the **Type of Binding** menu, select the type of binding for the pool:
   - **Unallocated**: No further configurations are required. Continue with the next step.
   - **Dynamic**: Configure the following settings:
     - **Network Address**. The subnet IP address for the dynamic DHCP pool.
     - **Network Mask**: The network mask for the dynamic DHCP pool.
     - **Network Prefix Length**: The subnet number for the dynamic DHCP pool. The range is from 0 to 32.

     NOTE: For a dynamic DHCP pool, you can configure either the network mask or the prefix length.

   - **Manual**: Configure the following settings:
     - **Client Name**: The DHCP client name.
     - **Hardware Address**: The hardware MAC address of the DHCP client.
     - **Hardware Address Type**: The protocol of the MAC address of the DHCP client. The type can be Ethernet or IEEE802. The default is Ethernet.
     - **Client ID**: The client ID of the DHCP client.
     - **Host Number**: The IP address for the manual binding of the DHCP client.
     - **Host Mask**: The subnet mask for the manual binding of the DHCP client.
     - **Host Prefix Length**: The subnet mask for the manual binding of the DHCP client. The range is from 0 to 32.

     NOTE: For a manually added DHCP client, you can configure either the network mask or the prefix length.

6. From the **Lease Time** menu, select how the binding is leased:
   - **Infinite**: For a dynamic binding, infinite is a lease period of 60 days. For manually added binding, the lease period is indefinite. No further configurations are required. Continue with the next step.
   - **Specified Duration**: Configure the following settings:
     - **Days**: The number of days of the lease period. The range is from 0 to 59. The default is 1.
     - **Hours**: The number of hours of the lease period. The range is from 0 to 22.
     - **Minutes**: The number of minutes of the lease period. The range is from 0 to 86399.
7. To configure default router addresses, click the **Default Router Addresses** link, and add up to eight IP addresses for default routers, in order of preference.
8. To configure DNS server addresses, click the **DNS Server Addresses** link, and add up to eight IP addresses for DNS servers, in order of preference.

9. To configure NetBIOS name server addresses, click the **NetBIOS name server addresses** link, and add up to eight IP addresses for NetBIOS name servers, in order of preference.
10. From the **NetBIOS Node Type** menu, select one of the following NetBIOS node types for DHCP clients:
    - **b-node Broadcast**
    - **p-node Peer-to-Peer**
    - **m-node Mixed**
    - **h-node Hybrid**
11. In the **Next Server Address** field, enter the IP address of the next server that must be used in the boot process for a DHCP client.
12. In the **Domain Name** field, enter the domain name that must be used in the boot process for a DHCP client.
    The domain name can be up to 255 characters in length.
13. In the **Bootfile** field, enter the name of the default boot image that must be used in the boot process for a DHCP client.
    The bootfile name can be up to 128 characters in length.
14. To configure NTP server addresses that must be used in the boot process for a DHCP client, click the **NTP Server** link, and add up to two IP address for NTP servers, in order of preference.
15. Click the **Add** button.
    The pool configuration is added.
16. To save the settings to the running configuration, click the **Save** icon.

Change a DHCP pool

You can change an existing DHCP pool.

To change a DHCP pool:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPServer >DHCP Pool Configuration**.
   The DHCP Pool Configuration page displays.
3. From the **Pool Name** menu, select the pool.
4. Change the settings as needed.
   For more information about the settings, see Create a DHCP pool on page 106.

5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Remove a DHCP pool

You can remove a DHCP pool that you no longer need.

To remove a DHCP pool:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Services > DHCP Server > DHCP Pool Configuration**.
   The DHCP Pool Configuration page displays.
3. From the **Pool Name** menu, select the pool.
4. Click the **Delete** button.
   The pool is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure DHCP pool options

You can specify information to include in the “options” field of a DHCP message that the switch sends.

To configure DHCP pool options:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPServer>DHCP Pool Options**.
   The DHCP Pool Options page displays.
3. From the **Pool Name** menu, select the pool name.
4. In the **Option Code** field, specify the option code configured for the pool.
   The range is from 1 to 254. For information about options and their format, see information about RFC 2132 at https://tools.ietf.org/html/rfc2132.

5. From the **Option Type** menu, select one of the following option types for the pool:
   - **ASCII**: Enter the information in the **Option Value** field in ASCII format.
   - **Hex**: Enter the information in the **Option Value** field in hexadecimal format. The possible formats are xx:xx and xxxx.xxxx.
   - **IP Address**: In the **Option Value** field, enter an IPv4 address.
6. In the **Option Value** field, specify the option value for the pool according to your selection in the previous step.
7. Click the **Add** button.
   The option code is added for the pool.
8. To save the settings to the running configuration, click the **Save** icon.

Bind a DHCP pool to a class and configure the binding

If you set up DHCP pools and classes for use with a DHCP relay (see DHCP relay on page 115), you can bind a pool to one or more classes and configure the bindings.
A binding configuration specifies the IP address ranges to which the binding applies, whether a DHCP address request is accepted or denied, and, as an option, the file path to a boot file that a client must download from a TFTP server in the network (this is referred to as Option Code 67).

To bind a DHCP pool to a class and configure the binding:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPServer >DHCP Class Binding Configuration**.
   The DHCP Class Binding Configuration page displays. The page shows different sections.
3. From the **Pool Name** menu, select the pool.
   The following steps apply to the DHCP Class Binding Configuration section.
4. From the **Class Name** field, select the class.
5. In the **IP Range From** field and **IP Range To** field, enter the start and end IP addresses to which the binding applies.
   If the binding must apply to all IP addresses in the class, enter 0.0.0.0 in each field.
6. From the **Deny Lease** menu, select if the DHCP request is accepted:
   - **FALSE**: The DHCP request is accepted. The devices that are defined by the IP address ranges can receive IP addresses from the DHCP server.
   - **TRUE**: The DHCP request is denied. The devices that are defined by the IP address ranges cannot receive IP addresses from the DHCP server.

7. Optionally, in the **Bootfile** field, enter the file path to a boot file on a TFTP server in the network.
   If the selection from the **Deny Lease** menu is **FALSE**, the devices download and install the bootfile from the TFTP server. The file path can be up to 128 characters.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Display DHCP server statistics

To display the DHCP server statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPServer >DHCP Server Statistics**.
   The DHCP Server Statistics page displays.
3. To refresh the page, click the **Refresh** button.
4. To clear the statistics, click the **Clear** button.
5. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 39.DHCP server statistics information

Field	Description
Binding Details	
Automatic Bindings	The number of automatic bindings on the DHCP server
Expired Bindings	The number of expired bindings on the DHCP server
Malformed Messages	The number of the malformed messages on the DHCP server
Message Received	
DHCPDISCOVER	The number of DHCPDISCOVER messages received by the DHCP server
DHCPREQUEST	The number of DHCPREQUEST messages received by the DHCP server
DHCPDECLINE	The number of DHCPDECLINE messages received by the DHCP server
DHCPRELEASE	The number of DHCPRELEASE messages received by the DHCP server
DHCPINFORM	The number of DHCPINFORM messages received by the DHCP server
Message Sent	
DHCPOFFER	The number of DHCPOFFER messages sent by the DHCP server
DHCPACK	The number of DHCPACK messages sent by the DHCP server
DHCPNAK	The number of DHCPNAK messages sent by the DHCP server
DHCP DISCOVER packets denied lease	The number of DHCP DISCOVER message that were denied by the DHCP server

Display the DHCP bindings

To display the DHCP bindings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **System>Services>DHCPServer >DHCP Bindings Information**.
   The DHCP Bindings Information page displays. The table displays information about all bindings.
3. To search for a binding by IP address, enter the IP address in the **Search Binding IP** field and click the **Go** button.
   Information about the binding displays.
   The following table describes the view-only fields on the page.

Table 40.DHCP bindings information

Field	Description
IP Address	The client's IP address
Pool Name	The name of the pool that the client is part of
Hardware Address	The client's hardware address
Lease Time Left	The remaining lease time in days, hours and minutes (dd:hh:mm format)
Type	The type of binding (Dynamic or Manual)

Delete one or all dynamic DHCP bindings

To delete one or all dynamic DHCP bindings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPServer >DHCP Bindings Information**.
   The DHCP Bindings Information page displays.
3. Select one of the following radio buttons:
   - **All Dynamic Bindings**: All dynamic bindings must be cleared.
   - **Specific Dynamic Binding**: In the field, specify the IP address of the dynamic binding that must be cleared.
4. Click the **Clear** button.
   The selected binding or bindings are deleted.
5. To save the settings to the running configuration, click the **Save** icon.

View bindings with DHCP conflicts

You can view information about devices with IP address conflicts. A conflict might occur when the same IP address is assigned to two or more devices in the network.

To view DHCP bindings with conflicts:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Services > DHCP Server > DHCP Conflicts Information**.
   The DHCP Conflicts Information page displays. The table displays information about all bindings with conflicts.
3. To search for a binding with an IP address conflict, enter the IP address in the **Search Conflict IP Address** field and click the **Go** button.
   Information about the binding displays.
   The following table describes the view-only fields on the page.

Table 41.DHCP conflicts information

Field	Description
IP Address	The IP address of the DHCP client.
Hardware Address	The hardware address if the DHCP client.
Detection Method	The method with which the IP address of the DHCP client was detected on the DHCP server.
Detection Time	The time when the conflict was detected in the days;hours;minutes;seconds format since the switch was last restarted.

Delete one or all DHCP bindings with conflicts

To delete one or all DHCP bindings with conflicts:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPServer >DHCP Conflicts Information**.
   The DHCP Conflicts Information page displays.

3. Select one of the following radio buttons:
   - **All Address Conflicts**: All bindings with IP address conflicts must be cleared.
   - **Specific Address Conflict**: In the field, specify the IP address of the binding with a conflict that must be cleared.
4. Click the **Clear** button.
   The selected binding or bindings are deleted.
5. To save the settings to the running configuration, click the **Save** icon.

DHCP relay

The switch can function as a Layer 3 device in which it relays messages between DHCP clients and DHCP servers that are located in a different IP subnet. We refer to the switch as a Layer 3 DHCP relay, or simply a DHCP relay, as opposed to a L2 DHCP relay, which is a different function (see DHCP Layer 2 relay on page 119).

Configure the global DHCP relay settings and display the relay statistics

You can enable the Layer 3 DHCP relay agent on the switch and let the switch relay DHCP messages between DHCP clients and DHCP servers that are located in a different IP subnet.

To configure the global DHCP relay settings and display relay statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services >DHCP Relay**.
   The DHCP Relay page display.
3. in the **Maximum Hop Count** field, enter the maximum number of hops a client request can take before being discarded.
   The range is from 1 to 16. The default is 4.
4. Select the Admin mode **Disable** or **Enable** radio button to disable or enable the DHCP relay on the switch.
5. In the **Minimum Wait Time** field, enter the minimum time in seconds that the DHCP relay must wait before forwarding requests.
   The time is compared to the time stamp in the client's DHCP request packets, which represents the time since the client was powered up. Packets are forwarded only when the time stamp exceeds the minimum wait time. The range is from 0 to 100.

6. Select one of the following Circuit ID Option Mode radio buttons:
   - **Disable**: Relay agent options are not added to requests before they are forwarded to the server and not removed from replies before they are forwarded to clients.
   - **Enable**: Relay agent options are added to requests before they are forwarded to the server and removed from replies before they are forwarded to clients.
   - **None**: The Circuit ID option is not configured. This selection is effectively the same as the Disable selection.
7. Select one of the following Server Override Mode radio buttons:
   - **Disable**: Suboption 5 and suboption 11 are not added to Option 82 of the incoming DHCP packet. Suboption 5 specifies the IP subnet from which the DHCP server must assign IP addresses.
     Suboption 11 specifies that the server ID can be replaced with a new ID.
   - **Enable**: Suboption 5 and suboption 11 are added to Option 82 of an incoming DHCP packet.
8. From the **Source Interface** menu, select an interface or VLAN.
   This is the interface or VLAN that is connected to the DHCP servers that are located in a different IP subnet.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

11. To refresh the information on the page, click the **Refresh** button.
    The following table describes the view-only fields on the page.

Table 42.DHCP relay status information

Field	Description
DHCP client messages received	The number of DHCP requests received from the clients.
DHCP client messages relayed	The number of DHCP requests forwarded to the servers.
DHCP server messages received	The number of DHCP server messages received from the servers.
DHCP server messages relayed	The number of DHCP server messages forwarded from the servers.
UDP client messages received	The number of UDP requests received from the clients. These are message that are transmitted over UDP from a DHCP client (external to the switch) and that are detected by the DHCP relay on the switch.
UDP client messages relayed	The number of UDP requests forwarded to the servers. These are message that are relayed over UDP to a DHCP server (external to the switch) and that are detected by the DHCP relay on the switch.
DHCP message hop count exceeded max	The number of DHCP client messages received for which the hop count exceeds the configured maximum hop count. These messages are not relayed.
DHCP message with secs field below min	The number of DHCP client messages received for which the minimum wait time is shorter than the configured minimum value. These messages are not relayed.
DHCP message with giaddr set to local address	The number of DHCP client messages received for which the gateway address (giaddr) is already set to an IP address configured on one of the relay agent's own IP addresses. In this situation, another device might be attempting to spoof the relay agent's address. These messages are not relayed.
Packets with expired TTL	The number of packets received with a time to live of 0 seconds or 1 second. These packets are not relayed.
Packets that matched a discard entry	The number of packets ignored by the relay agent because the packets match a discard relay entry.

Configure a DHCP relay interface

You can enable the Layer 3 DHCP relay agent on a routing interface and let the interface relay DHCP messages between DHCP clients and DHCP servers that are located in a different IP subnet.

To configure a DHCP relay interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services >DHCP Relay**.
   The DHCP Relay page display.
3. Scroll down to the lower DHCP relay section at the bottom of the page.
   If no routing interfaces are configured on the switch, the DHCP Relay configuration table does not display on the page.
4. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
5. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

6. Select one of the following Server Override Mode radio buttons:
   - **Disable**: Suboption 5 and suboption 11 are not added to Option 82 of the incoming DHCP packet. Suboption 5 specifies the IP subnet from which the DHCP server must assign IP addresses.
     Suboption 11 specifies that the server ID can be replaced with a new ID.
   - **Enable**: Suboption 5 and suboption 11 are added to Option 82 of an incoming DHCP packet.
7. From the **Source Interface** menu, select an interface or VLAN.
   This is the interface or VLAN that is connected to the DHCP servers that are located in a different IP subnet.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

DHCP Layer 2 relay

If you enable and configure a DHCP relay agent on the switch, some Layer 2 (L2) devices do not need to connect to a DHCP server on the physical network. A relay agent automatically populates the gateway IP address (giaddr) field and adds the Relay Agent Information option to DHCP messages. A DHCP server uses this option for IP addresses and other assignment policies. A DHCP relay agent is usually an IP routing-aware device, which is also referred to as a Layer 3 relay agent. In some network configurations, L2 devices must append the Relay Agent Information option because they are closer to the end hosts.
An L2 device can operate as a bridge only for a network and might not include an IPv4 address on the network. If an L2 device lacks an IPv4 source address, the device cannot relay packets directly to a DHCP server that is located on another network. In that situation, the L2 device can append the Relay Agent Information option and broadcast the DHCP message.

Configure the global DHCP L2 relay settings

To configure the global DHCP L2 relay settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPL2Relay >DHCP L2 Relay Global Configuration**.
   The DHCP L2 Relay Global Configuration page displays.

3. Select the Admin mode **Disable** or **Enable** radio button to disable or enable the DHCP L2 relay on the switch.
   The default is Disable.
4. In the DHCP L2 Relay VLAN Configuration section, configure the following settings:
   - **VLAN ID**: The ID of the VLAN that is already configured on the switch. You cannot use this page to add a VLAN. For information about adding VLANs, see VLANs on page 195.
   - **Admin mode**: Enable or disable the DHCP L2 relay on the VLAN.
   - **Circuit ID mode**: Enable or disable the circuit ID suboption of DHCP Option-82.
   - **Remote ID String**: Specify the remote ID if the Remote ID option is enabled.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Configure a DHCP L2 relay interface

You can configure a DHCP L2 relay on an interface.

To configure DHCP L2 relay interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPL2Relay >DHCP L2 Relay Interface Configuration**.
   The DHCP L2 Relay Interface Configuration displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.

4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Admin mode** menu, select to enable or disable the DHCP L2 relay on the interface. The default is Disable.
6. From the **82 Option Trust Mode** menu, select to enable or disable the interface as a trusted interface for Relay Agent Information option (Option 82).
7. From the **No 82 Option Trust Mode** menu, select how the interface treats packets that are received without Option 82:
   - **Drop**: The interface drops packets that are received without Option 82.
   - **Update**: The interface adds Option 82 to packets that are received without Option 82.

   The settings of this menu can be in effect only if you select **Enable** from the **82 Option Trust Mode** menu.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Display DHCP L2 relay interface statistics

To display the DHCP L2 relay interface statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPL2Relay >DHCP L2 Relay Interface Statistics**.
   The DHCP L2 Relay Interface Statistics page displays.

3. Select whether to display physical interfaces, LAGs, VLANs, or all by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Physical interfaces, LAGs, and VLANs are displayed, or for a switch stack, physical interfaces on all switches in the stack, LAGs, and VLANs are displayed.
4. To refresh the information on the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 43.DHCP L2 relay interface statistics information

Field	Description
Interface	The interface from which the DHCP message is received
UntrustedServerMsgsWithOpt82	The number of DHCP messages with Option 82 received from an untrusted server
UntrustedClientMsgsWithOpt82	The number of DHCP messages with Option 82 received from an untrusted client
TrustedServerMsgsWithoutOpt82	The number of DHCP messages without Option 82 received from a trusted server
TrustedClientMsgsWithoutOpt82	The number of DHCP messages without Option 82 received from a trusted client

UDP relay

You can enable the switch to function as a UDP relay, which enables the switch to forward User Datagram Protocol (UDP) broadcast packets from a router to an IP address on a local or non-local subnet.

Configure the global UDP relay settings and add a UDP relay

To configure the global UDP relay settings and add a UDP relay:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **System>Services>UDPRelay >UDP Relay Global Configuration**.
   The UDP Relay Global Configuration page displays.
3. Select the Admin mode **Disable** or **Enable** radio button to disable or enable the UDP relay on the switch.
   The default is Disable.
4. In the UDP Relay Global Configuration section, configure the following settings:
   - **Server Address**: Enter the IPv4 address of the UDP relay server.
   - **UDP Port**: Select one of the following types of UDP destination ports:
     - **Other**: The **UDP Port Other Value** field becomes available so that you can enter a custom UDP port number.
     - **dhcp**: Relay DHCP (UDP port 67) packets.
     - **domain**: Relay DNS (UDP port 53) packets.
     - **isakmp**: Relay ISAKMP (UDP port 500) packets.
     - **mobile-ip**: Relay Mobile IP (UDP port 434) packets
     - **nameserver**: Relay IEN-116 Name Service (UDP port 42) packets
     - **netbios-dgm**: Relay NetBIOS Datagram Server (UDP port 138) packets
     - **netbios-ns**: Relay NetBIOS Name Server (UDP port 137) packets
     - **ntp**: Relay network time protocol (UDP port 123) packets.
     - **pim-auto-rp**: Relay PIM auto RP (UDP port 496) packets.
     - **rip**: Relay Routing Image Protocol (RIP) (UDP port 520) packets
     - **tacacs**: Relay TACACS (UDP port 49) packet
     - **tftp**: Relay TFTP (UDP port 69) packets
     - **time**: Relay time service (UDP port 37) packets
   - **UDP Port Other Value**: If you select **Other** from the **UDP Port** menu, enter a custom UDP port number in the range from 0 to 65535.
5. Click the **Add** button.
   The UDP relay is added.
   The Hit Count field displays the number of UDP packets that are detected on the UDP port for the UDP relay.
6. To save the settings to the running configuration, click the **Save** icon.

Change a UDP relay configuration

To change a UDP relay configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **System>Services>UDPRelay >UDP Relay Global Configuration**.
   The UDP Relay Global Configuration page displays.
3. In the UDP Relay Global Configuration section, select the check box for the UDP relay configuration.
4. Change the settings as needed.
   For more information about the settings, see Configure the global UDP relay settings and add a UDP relay on page 122.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Remove a UDP relay configuration

To remove a UDP relay configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Services > UDP Relay > UDP Relay Global Configuration**.
   The UDP Relay Global Configuration page displays.
3. In the UDP Relay Global Configuration section, select the check box for the UDP relay configuration.
4. Click the **Delete** button.
   The UDP switch configuration is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Add a UDP interface configuration

To add a UDP relay interface configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>UDPRelay >UDP Relay Interface Configuration**.
   The page UDP Relay Interface Configuration page displays.
3. From the **Interface** menu, select the interface or VLAN for which you are adding the UDP configuration.

4. in the **Server Address** field, enter the IPv4 address of the UDP relay server.
5. From the **UDP Port** menu, select one of the following types of UDP destination ports:
   - **Other**: The **UDP Port Other Value** field becomes available so that you can enter a custom UDP port number.
   - **dhcp**: Relay DHCP (UDP port 67) packets.
   - **domain**: Relay DNS (UDP port 53) packets.
   - **isakmp**: Relay ISAKMP (UDP port 500) packets.
   - **mobile-ip**: Relay Mobile IP (UDP port 434) packets
   - **nameserver**: Relay IEN-116 Name Service (UDP port 42) packets
   - **netbios-dgm**: Relay NetBIOS Datagram Server (UDP port 138) packets
   - **netbios-ns**: Relay NetBIOS Name Server (UDP port 137) packets
   - **ntp**: Relay network time protocol (UDP port 123) packets.
   - **pim-auto-rp**: Relay PIM auto RP (UDP port 496) packets.
   - **rip**: Relay Routing Image Protocol (RIP) (UDP port 520) packets
   - **tacacs**: Relay TACACS (UDP port 49) packet
   - **tftp**: Relay TFTP (UDP port 69) packets
   - **time**: Relay time service (UDP port 37) packets
6. If you select **Other** from the **UDP Port** menu, enter a custom UDP port number in the **UDP Port Other Value**. field.
   The port number can be in the range from 0 to 65535.
7. From the **Discard** menu, select if packets that match the condition are dropped:
   - **Enable**: The packets are dropped. This selection functions only is you enter 0.0.0.0 as the server IP address.
   - **Discard**: The packets are not dropped. This selection functions if the server IP address is not 0.0.0.0.
8. Click the **Add** button.
   The UDP interface configuration is added.
   The Hit Count field displays the number of UDP packets that are detected on the UDP interface.
9. To save the settings to the running configuration, click the **Save** icon.

Change a UDP interface configuration

To change a UDP switch configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **System>Services>UDPRelay >UDP Relay Interface Configuration**.
   The page UDP Relay Interface Configuration page displays.
3. From the **Interface** menu, select the interface or VLAN.
4. Change the settings as needed.
   For more information about the settings, see Add a UDP interface configuration on page 124.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Remove a UDP interface

To remove a UDP interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Services > UDP Relay > UDP Relay Interface Configuration**.
   The page UDP Relay Interface Configuration page displays.
3. From the **Interface** menu, select the interface or VLAN.
4. Click the **Delete** button.
   The UDP interface is removed.
5. To save the settings to the running configuration, click the **Save** icon.

DHCPv6 server

You can configure settings for a DHCPv6 server, DHCPv6 pools, DHCPv6 prefix delegation, DHCPv6 interfaces. and DHCPv6 bindings. You can also view DHCPv6 statistics.

Enable the DHCPv6 server

By default, the DHCPv6 server is disabled. You can enable it so that the switch can assign network configuration information to IPv6 clients.

To enable the DHCPv6 server on the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **System > Services > DHCPv6 Server > DHCPv6 Server Configuration**.
   The DHCPv6 Server Configuration page displays.
3. Select the Admin mode **Enable** radio button.
   The default value is Disable.
4. Click the **Apply** button.
   Your settings are saved.
   The DHCPv6 Server DUID field displays the DHCP Unique Identifier (DUID) of the DHCPv6 server.
5. To save the settings to the running configuration, click the **Save** icon.

Manage DHCPv6 pools

A DHCPv6 pool provides network configuration information that is available to DHCPv6 clients that request such information. You can add, change, and remove DHCPv6 pools.

Create a DHCPv6 pool

A DHCPv6 pool provides network configuration information that is available to DHCPv6 clients that request such information.

To create a DHCPv6 pool:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPv6Server >DHCPv6 Pool Configuration**.
   The DHCPv6 Pool Configuration page displays.
3. From the **Pool Name** menu, select **Create**.
4. In the **Pool Name** field, type a name that identifies the DHCPv6 server pool.
   The name can be up to 31 characters in length.
5. Click the **DNS Server Addresses** link and specify the list of default IPv6 router addresses for the pool.
   You can specify up to eight default router addresses in order of preference.
6. Click the **Domain Name** link and specify the list of IPv6 domain names for the pool.
   A domain name can be up to 255 alphanumeric characters in length.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Change a DHCPv6 pool

You can change an existing DHCPv6 pool.

To change a DHCPv6 pool:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPv6Server >DHCPv6 Pool Configuration**.
   The DHCPv6 Pool Configuration page displays.
3. From the **Pool Name** menu, select the name for the pool.
   You cannot change the pool name.
4. Click the **DNS Server Addresses** link and change the list of default IPv6 router addresses for the pool.
   You can specify up to eight default router addresses in order of preference.
5. Click the **Domain Name** link and change the list of IPv6 domain names for the pool.
   A domain name can be up to 255 alphanumeric characters in length.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Delete a DHCPv6 pool

You can delete a DHCPv6 pool that you no longer need.

To delete a DHCPv6 pool:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPv6Server >DHCPv6 Pool Configuration**.
   The DHCPv6 Pool Configuration page displays.
3. From the **Pool Name** menu, select the name for the pool.
4. Click the **Delete** button.
   The pool is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Manage DHCPv6 prefix delegation for pools

A prefix delegation configuration supplies an IPv6 prefix and prefix length to a DHCPv6 pool that provide network configuration information to DHCPv6 clients that request such information. A DHCPv6 client can request multiple IPv6 prefixes and can also request specific IPv6 prefixes. If a DHCPv6 pool contains the specific prefix that a DHCPv6 client requests, the prefix is delegated to the client. Otherwise, the first available IPv6 prefix in a pool is delegated to the client.

Create a DHCPv6 prefix delegation configuration for a pool

You can create a DHCPv6 prefix delegation configuration for a pool.

To create a DHCPv6 prefix delegation configuration for a pool:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPv6Server >DHCPv6 Prefix Delegation Configuration**.
   The DHCPv6 Prefix Delegation Configuration page displays.
3. From the **Pool Name** menu, select a pool.
4. In the **Prefix** field, specify the IPv6 prefix.
5. In the **Prefix Length** field, specify the length that is associated with the IPv6 prefix.
6. In the **DUID** field, specify the DHCP Unique Identifier (DUID) that is associated with the IPv6 prefix.

   NOTE: The DUID and client name are used only if a device requests a prefix. In this way, the device receives an IPv6 prefix instead of specific IPv6 address. If you assign a specific IPv6 address to the device, you can leave the DUID and Client Name fields blank.

7. In the **Client Name**, field, enter a client name that is associated with the IPv6 prefix.
   The name, which is useful for logging or tracing, can be up to 31 characters.
8. In the **Valid Lifetime** field, enter the time in seconds after which the IPv6 prefix times out.
   Enter a value between 0 and 4294967295.
9. In the **Prefer Lifetime** field, enter the preferred time in seconds after which the IPv6 prefix times out.
   Enter a value between 0 and 4294967295.

10. Click the **Add** button.
    The prefix delegation configuration is added for the selected pool.
11. To save the settings to the running configuration, click the **Save** icon.

Change a DHCPv6 prefix delegation configuration for a pool

You can change an existing DHCPv6 prefix delegation configuration.

To change a DHCPv6 prefix delegation configuration for a pool:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPv6Server >DHCPv6 Prefix Delegation Configuration**.
   The DHCPv6 Prefix Delegation Configuration page displays.
3. From the **Pool Name** menu, select the pool.
4. Change the settings as needed.
   For more information about the settings, see Create a DHCPv6 prefix delegation configuration for a pool on page 129.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete a DHCPv6 prefix delegation configuration for a pool

You can delete an existing DHCPv6 prefix delegation configuration that you no longer need.

To delete a DHCPv6 prefix delegation configuration for a pool:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPv6Server >DHCPv6 Prefix Delegation Configuration**.
   The DHCPv6 Prefix Delegation Configuration page displays.
3. From the **Pool Name** menu, select the pool.
4. Click the **Delete** button.
   The prefix delegation configuration is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the DHCPv6 settings for an interface

You can configure the DHCPv6 settings for an interface so that the interface can provide DHCPv6 services to attached devices.

To configure the DHCPv6 settings for an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPv6Server >DHCPv6 Interface Configuration**.
   The DHCPv6 Interface Configuration page displays.
3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Admin mode** menu, select to enable or disable the DHCPv6 server mode on the interface.

   NOTE: DHCPv6 server and DHCPv6 relay functions are mutually exclusive. For information about configuring an interface as a DHCPv6 relay, see DHCPv6 relay interface on page 136.
6. From the **Pool Name** field, select a DHCPv6 pool.
   For more information, see Manage DHCPv6 pools on page 127.

7. From the **Rapid Commit** menu, select to enable or disable an abbreviated exchange between the device that requests IPv6 information and the DHCPv6 server.
   This setting is optional.
8. In the **Preference** field, specify the preference value that a device can use to determine the preference of this interface in relation to other DHCPv6 servers.
   You can enter a value from 0 to 4294967295. The default is 0.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

Display the DHCPv6 bindings

After a client acquires IPv6 configuration information from a DHCPv6 server, the server adds an entry to the DHCPv6 bindings table.

To display DHCPv6 bindings information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPv6Server >DHCPv6 Bindings Information**.
   The DHCPv6 Bindings Information page displays.
3. To search for a binding, enter the IPv6 address in the **Search Binding IP** field and click the **Go** button.
   Information about the binding displays.
4. The following table describes the view-only fields on the page.

5. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 44.DHCPv6 binding Information

Field	Description
Client Address	The IPv6 address of the client associated with the binding.
Client Interface	The interface number on which the client binding occurred.
Client DUID	The DHCPv6 Unique Identifier (DUID) of the client. The DUID is a combination of the client's hardware address and client identifier.
Prefix	The IPv6 address for the delegated prefix associated with this binding.
Prefix Length	The IPv6 mask length for the delegated prefix associated with this binding.
Prefix Type	The type of IPv6 prefix associated with this binding.
Expiry Time	The number of seconds until the prefix associated with a binding expires.
Valid Lifetime	The maximum amount of time in seconds that the client is allowed to use the prefix.
Prefer Lifetime	The preferred amount of time in seconds that the client is allowed to use the prefix.

Display DHCPv6 server statistics

You can display the DHCPv6 server statistics, including information about the DHCPv6 messages, sent, received, and discarded globally and on each interface.

To display DHCPv6 server statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPv6Server >DHCPv6 Server Statistics**.
   The DHCPv6 Server Statistics page displays.
3. Do one of the following:
   - **Display statistics for one interface**: From the Interface menu, select the interface.
   - **Display statistics for all interfaces**: From the Interface menu, select **All**.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 45.DHCPv6 server statistics information

Field	Description
Messages Received	
Total DHCPv6 Packets Received	The number of DHCPv6 messages received. The DHCPv6 messages sent from a DHCP v6 client to a DHCP v6 server include solicit, request, confirm, renew, rebind, release, decline, and information-request messages. Additionally, a DHCP v6 relay agent can forward relay-forward messages to a DHCP v6 server.
DHCPv6 Solicit Packets Received	The number of DHCPv6 Solicit messages received. This type of message is sent by a client to locate DHCPv6 servers.
DHCPv6 Request Packets Received	The number of requests received.
DHCPv6 Confirm Packets Received	The number of DHCPv6 Confirm messages received. This type of message is sent by a client to all DHCPv6 servers to determine whether its configuration is valid for the connected link.
DHCPv6 Renew Packets Received	The number of DHCPv6 Renew messages received. This type of message is sent by a client to extend and update the configuration information provided by the DHCPv6 server.
DHCPv6 Rebind Packets Received	The number of DHCPv6 Rebind messages received. This type of message is sent by a client to any DHCPv6 server when it does not receive a response to a Renew message.
DHCPv6 Release Packets Received	The number of DHCPv6 Release messages received. This type of message is sent by a client to indicate that it no longer needs the assigned address.
DHCPv6 Decline Packets Received	The number of DHCPv6 Decline messages received. This type of message is sent by a client to the DHCPv6 server to indicate that an assigned address is already in use on the link.
DHCPv6 Inform Packets Received	The number of DHCP v6 information-request messages received. This type of message is sent by a client to request configuration information other than IP address assignment.
DHCPv6 Relay-forward Packets Received	The number of DHCPv6 relay-forward messages received. This type of message is sent by a relay agent to forward messages to

Field	Description
	servers.
DHCPv6 Relay-reply Packets Received	The number of DHCP v6 relay-reply messages received. This type of message is sent by a server to a DHCP v6 relay agent and contains the message for the relay agent to deliver to the client.
DHCPv6 Malformed Packets Received	The number of DHCPv6 messages received but dropped because they were malformed.
Received DHCPv6 Packets Discarded	The number of packets discarded.
Messages Sent	
Total DHCPv6 Packets Sent	The number of DHCPv6 messages sent. The DHCPv6 messages sent from a DHCPv6 server to a DHCPv6 client include Advertise, Reply, Reconfigure, and Relay-Reply messages.
DHCPv6 Advertisement Packets Transmitted	The number of DHCPv6 Advertise messages sent. This type of message is sent by a server to a DHCPv6 client in response to a Solicit message and indicates that it is available for service.
DHCPv6 Reply Packets Transmitted	The number of DHCPv6 Reply messages sent to a DHCPv6 client in response to a solicit, request, renew, rebind, information-request, confirm, release, or decline message.
DHCPv6 Reconfig Packets Transmitted	The number of DHCPv6 reconfigure messages sent. This type of message is sent by a server to a DHCPv6 client to inform the client that the server has new or updated information. The client then typically initiates a renew/reply or Information-request/ reply transaction with the server to receive the updated information.
DHCPv6 Relay-forward Packets Transmitted	The number of DHCPv6 Relay-Forward messages sent. This type of message is sent by a relay agent to forward messages to servers.
DHCPv6 Relay-reply Packets Transmitted	The number of DHCPv6 Relay-Reply messages sent. This type of message is sent by a server to a DHCPv6 relay agent and contains the message for the relay agent to deliver to the client.

Delete DHCPv6 statistics for one or all interfaces

To delete DHCPv6 statistics for one or all interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>Services>DHCPv6Server >DHCPv6 Server Statistics**.
   The DHCPv6 Server Statistics page displays.
3. Do one of the following:
   - **Delete statistics for one interface**: From the Interface menu, select the interface.
   - **Delete statistics for all interfaces**: From the Interface menu, select **All**.
4. Click the **Clear** button.
   The selected statistics are deleted.
5. To save the settings to the running configuration, click the **Save** icon.

DHCPv6 relay interface

A DHCPv6 relay agent allows sub-options to be attached to messages that are being relayed by the switch to a DHCPv6 server. In turn, the DHCPv6 server can use this information in determining an address to assign to a DHCPv6 client.

To configure an interface as a DHCPv6 relay:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Services > DHCPv6 Relay**.
   The DHCPv6 Interface Configuration page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

5. From the **Admin mode** menu, select to enable or disable the DHCPv6 relay mode on the interface.

> **NOTE:** DHCPv6 relay and DHCPv6 server functions are mutually exclusive. For information about configuring an interface as a DHCPv6 server, see Configure the DHCPv6 settings for an interface on page 131.

6. From the **Relay Interface** menu, select the interface, LAG, or VLAN that is connected to the DHCPv6 relay server.
7. In the **Destination IP Address** field, specify the IPv6 address for the DHCPv6 relay server.
8. In the **Remote ID** field, specify the relay agent information option.
   The remote ID is derived from the DHCPv6 server DUID and the relay interface number, or you can specify is as a user-defined string.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

Power over Ethernet

You can configure the global Power over Ethernet (PoE) configuration settings and the PoE settings for each port.

PoE concepts

The Power over Ethernet (PoE) models support PoE+ or PoE++ ports with the port capacities and budgets that are described in the following tables.

Table 46.PoE port capacities

Model	PoE ports	Port Capacity
M4350-24X4V	24 PoE+ (802.3at)	30W
M4350-24G4XF	24 PoE+ (802.3at)	30W
M4350-48G4XF	48 PoE+ (802.3at)	30W
M4350-44M4X4V	48 PoE++ (802.3bt)	90W
M4350-36X4V	36 PoE++ (802.3bt)	90W
M4350-24X8F8V	24 PoE++ (802.3bt)	90W
M4350-40X4C	40 PoE++ (802.3bt)	90W

Model	PoE ports	Port Capacity
M4350-8M2V	8 PoE++ (802.3bt)	90W
M4350-24M4X4V	28 PoE++ (802.3bt)	90W
M4350-16M4V	16 PoE++ (802.3bt)	90W

Table 47.PoE switch budgets

Model	Switch PoE Budget	
M4350-24X4V (1 PSU bay for optional APS)	Internal PSU only	576W
	Add 1 x APS350W	700W
	Add 1 x APS600Wv21	720W
M4350-24G4XF (1 PSU bay for optional APS)	Internal PSU only	648W
	Add 1 x APS350W2	720W
M4350-48G4XF (2 PSU bays for optional APS)	Internal PSU only	236W
	Add 1 x APS350W	436W
	Add 1 x APS600Wv2	636W
	Add 1x APS920W	892W
	Add 1x APS2000W at 110V	956W
	Add 1x APS2000W at 220V	1440W
	Add 2 x APS350W	716W
	Add 2 x APS600Wv2	1116W
	Add 2x APS920W	1440W
	Add 2x APS2000W at 110V	1440W
	Add 2x APS2000W at 220V	1440W
M4350-44M4X4V (2 PSU bays for optional APS)	Internal PSU only	194W
	Add 1 x APS350W	394W
	Add 1 x APS600Wv2	594W
	Add 1x APS920W	850W
	Add 1x APS2000W at 110V	914W
	Add 1x APS2000W at 220V	1786W

Model	Switch PoE Budget	
	Add 2 x APS350W	674W
	Add 2 x APS600Wv2	1074W
	Add 2x APS920W	1586W
	Add 2x APS2000W at 110V	1714W
	Add 2x APS2000W at 220V	3314W
M4350-36X4V (1 PSU bay for optional APS)	Internal PSU only	280W
	Add 1 x APS600Wv3	640W
	Add 1 x APS1200Wv2 at 110V	960W
	Add 1 x APS1200Wv2 at 220V	1120W
	Add 1 x APS2000Wv2 at 110V	960W
	Add 1 x APS2000Wv2 at 220V	1760W
M4350-24X8F8V (1 PSU bay for optional APS)	Internal PSU only	290W
	Add 1 x APS600Wv3	650W
	Add 1 x APS1200Wv2 at 110V	970W
	Add 1 x APS1200Wv2 at 220V	1130W
	Add 1 x APS2000Wv2 at 110V	970W
	Add 1 x APS2000Wv2 at 220V	1770W
M4350-40X4C (1 PSU bay for optional APS)	Internal PSU only	196W
	Add 1 x APS600Wv3	556W
	Add 1 x APS1200Wv2 at 110V	876W
	Add 1 x APS1200Wv2 at 220V	1036W
	Add 1 x APS2000Wv2 at 110V	876W
	Add 1 x APS2000Wv2 at 220V	1676W
M4350-8M2V (3 AC power receptacles for optional APS)	Internal PSU only	119W
	Add 2 x APS254W	551W
M4350-24M4X4V (2 PSU bays for optional APS)	Internal PSU only	552W
	Add 1 x APS600Wv3	1332W
	Add 2 x APS1200Wv2	1972W

Model	Switch PoE Budget	
	Add 1 x APS2000Wv2	2520W
M4350-16M4V	Internal PSU 1 only	530W
	Add internal PSU 2	1130W

1. If you install a more powerful APS such as an APS920W or APS2000W, the PoE budget does not increase.

2. If you install a more powerful APS such as an APS600Wv2, APS920W, or APS2000W, the PoE budget does not increase.

Supplied power is prioritized according to the port order, up to the total power budget of the device. The lowest-numbered PoE port (for example, port 1) receives the highest PoE priority, while the highest-numbered PoE port (for example, port 24) is relegated to the lowest PoE priority.

If the power requirements for attached powered devices (PDs) exceed the total power budget of the switch, the PoE power to the device on the highest-numbered active PoE port is disabled to make sure that the devices connected to the higher-priority, lower-numbered PoE ports are supported first.

Although a device might be listed as an 802.3bt PoE++-powered or 802.3at PoE+-powered device, it might not require the maximum power limit that is specified by its IEEE standard. Many devices require less power, allowing all 8 PoE ports to be active simultaneously when the devices correctly report their PoE class to the switch.

The following table shows the standard power ranges, calculated with the maximum cable length of 328 feet (100 meters). If a device receives insufficient PoE power from the switch, consider using a shorter cable.

Table 48.PoE classes and PoE power allocations

Device Class	Compatible PoE Standard	Class Description	Maximum Power Reserved for the PD	Power Delivered to the PD
0	PoE, PoE+, and PoE++	Default power (full)	15.4W	0.44W–15.8W
1	PoE, PoE+, and PoE++	Very low power	4.0W	0.44W–3.84W
2	PoE, PoE+, and PoE++	Low power	7.0W	3.84W–7.2W
3	PoE, PoE+, and PoE++	Mid power	15.4W	6.49W–15.9W
4	PoE+ and PoE++	High power	30.0W	12.95W–30.8W
5	PoE++	Ultra high power	45.0W	25.5W–47.0W
6	PoE++	Ultra high power	60.0W	51.0W–64.4W

Device Class	Compatible PoE Standard	Class Description	Maximum Power Reserved for the PD	Power Delivered to the PD
7	PoE++	Ultra high power	75.0W	62.0W–81.1W
8	PoE++	Ultra high power	90.0W	71.0W–96.5W

Set the PoE system usage threshold and power management mode

You can configure a threshold for the PoE usage level at which a trap is sent and you can set the power management mode that the switch uses to deliver power to the requesting powered devices (PDs).

To set the PoE system usage threshold and power management mode:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>PoE >Basic >PoE Configuration**.
   The PoE Configuration page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.
4. In the **System Usage Threshold** field, enter a number from 1 to 99 to set the threshold level at which a trap is sent if the consumed power exceeds the threshold power.
5. From the **Power Management mode** menu, select the power management algorithm that the switch uses to deliver power to the requesting PDs:
   - **Static**: Select **Static** to specify that the power allocated for each port depends on the type of power threshold that is configured on the port.
   - **Dynamic**: Select **Dynamic** to specify that the power consumption on each port is measured and calculated in real time. This is the default setting.
6. To set the traps, in the PoE Trap Configuration section, select one of the following radio buttons:
   - **Enable**: Activates the PoE traps, which means that the switch can send PoE traps. This is the default setting.
   - **Disable**: Deactivates the PoE traps, which means that the switch cannot not send PoE traps.

7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 49.PoE information

Field	Description
Unit	The unit is always 1, unless you configured a switch stack, in which the case the unit depends on the stacked switch.
Slot	The slot is always 0.
Model	The power sourcing equipment (PSE) controller
Host	The switch model in which the PSE controller is installed
Firmware Version	The firmware version of the PSE controller
Power Status	The power status (Off or On)
Total Power Available (W)	The maximum power in watts the switch can deliver to all ports.
Threshold Power (W)	If the consumed power is below the threshold power, the switch can power up another port. The consumed power can be between the nominal and threshold power. The threshold power is displayed in watts. **Note**: The threshold power value is determined by the value that you enter in the System Usage Threshold field.
Consumed Power (W)	Total power in watts that is being delivered to all ports.

Configure the PoE ports settings

To configure the PoE ports settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System>PoE >Advanced >PoE Port Configuration**.
   The PoE Port Configuration page displays.

3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.
4. Select one or more ports by taking one of the following actions:
   - To configure a single port, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple ports with the same settings, select the check box associated with each port.
   - To configure all ports with the same settings, select the check box in the heading row.
5. From the **Port Power** menu, select the administrative PoE mode of the port:
   - **Enable**: The port's capacity to deliver power is enabled. This is the default setting.
   - **Disable**: The port's capacity to deliver power is disabled.
6. From the **Port Priority** menu, select the priority for the port in relation to other ports if the total power that the switch is capable of delivering exceeds the total power budget:
   - **Low**: Low priority. This is the default setting.
   - **Medium**: Medium priority.
   - **High**: High priority.
   - **Critical**: Critical priority.

   The port priority determines which ports can still deliver power after the total power delivered by the switch exceeds the total power budget. (In such a situation, the switch might not be able to deliver power to all connected devices.) If the same priority applies to two ports, the lower-numbered port receives higher priority.

7. From the **Power Mode** menu, select the PoE mode that the port must function in:
   - **802.3af**: The port is powered in and limited to the IEEE 802.3af mode. A PD that requires IEEE 802.3at does not receive power if the port functions in IEEE 802.3af mode.
   - **Legacy**: The port is powered using high-inrush current, which is used by legacy PDs that require more than 15W to power up.
   - **Pre-802.3at**: The port is initially powered in the IEEE 802.3af mode and, before 75 msec pass, is switched to the high power IEEE 802.3at mode. Select this mode if the PD does not perform Layer 2 classification or if the switch performs 2-event Layer 1 classification.
   - **802.3at**: The port is powered in the IEEE 802.3at mode and is backward compatible with IEEE 802.3af. The 802.3at mode is the default mode. In this mode, if the switch detects that the attached PD requests more power than IEEE 802.3af but is not an IEEE 802.3at Class 4 device, the PD does not receive power from the switch.
   - **Pre-802.3bt** (PoE++ models only): The port supports Class 4 devices that use 4-pair PoE (4PPoE) to receive power higher than 30W but that are not compliant with IEEE 802.3bt. The port also supports the IEEE 802.3at and IEEE 802.3af modes.
   - **802.3bt-Type3** (PoE++ models only): The port supports the IEEE 802.3bt Type 3 mode, the IEEE 802.3at mode, and the IEEE 802.3af mode.
   - **802.3bt** (PoE++ models only): The port is powered in the IEEE 802.3bt mode and is backward compatible with IEEE 802.3at and IEEE 802.3af. In this mode, if the switch detects that the attached PD requests more power than IEEE 802.3at but is not an IEEE 802.3bt device, the PD does not receive power from the switch.

8. From the **Power Limit Type** menu, select how the port controls the maximum power that it can deliver:
   - **None**: For PoE+ (802.3at) ports, the port draws up to Class 0 maximum power in low power mode. In high power mode, the following applies:
     - **PoE+ (802.3at) ports**: The port draws up to Class 4 maximum power.
     - **PoE++ (802.3bt) ports** (PoE++ models only): The port draws up to Class 8 maximum power.
   - **Class**: The port power limit is equal to the class of the attached PD. This is the default setting. The upper limit is the power that a port can deliver to a PD. The class is detected based on the PD that is attached to the port, and the following applies:
     - **PoE+ (802.3at) ports**: Possible values are from Class 0 to Class 4.
     - **PoE++ (802.3bt) ports** (PoE++ models only): Possible values are from Class 0 to Class 8.
   - **User**: The port power limit is equal to the value that is specified in the **Power Limit (W)** field.

   NOTE: If a PD does not report its class correctly, use of these options can preserve additional PoE power by preventing the switch from delivering more power than the PD requires. However, depending on which option you select, a PD that does not report its class correctly might not power up at all.

9. In the **Power Limit (W)** field, enter the maximum power (in W) that the port can deliver.
   The following applies:
   - **PoE+ (802.3at) ports**: The range is from 3.0W to 30.0W.
   - **802.3bt-Type3 ports** (which is a selection from the **Power Mode** menu for PoE++ models only): The range is from 3.0W to 60.0W.
   - **PoE++ (802.3bt) ports** (PoE++ models only): The range is from 3.0W to 99.9W.
10. From the **Detection Type** menu, select how the port detects the attached PD:
    - **IEEE 802**: The port performs a 4-point resistive detection. This is the default setting.
    - **4pt 802.3af + Legacy**: The port performs a 4-point resistive detection, and if required, continues with legacy detection.
    - **Legacy**: The port performs legacy detection.
11. From the **TimerSchedule** menu, select a timer schedule or select **None**, which is the default selection.
    For information about setting up and configuring PoE timer schedules, see Timer schedules on page 153.
12. Click the **Apply** button.
    Your settings are saved.

13. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 50.PoE port information

Field	Description
High Power	All ports supports high power mode.
Max Power (W)	The maximum power in Watts that can be provided by the port.
Class	The class defines the range of power that a powered device (PD) is drawing from the switch. The class definitions are as follows: • **0**: 0.44W–15.8W • **1**: 0.44W–3.84W • **2**: 3.84W–7.2W • **3**: 6.49W–15.9W • **4**: 12.95W–30.8W • **5**: 25.5W–47.0W (PoE++ models only) • **6**: 51.0W–64.4W (PoE++ models only) • **7**: 62.0W–81.1W (PoE++ models only) • **8**: 71.0W–96.5W (PoE++ models only) • **Unknown**: The class cannot be detected, or no PD is attached to the port.
Output Voltage (Volts)	The voltage that is delivered to the PD in volts.
Output Current (mA)	The current that is delivered to the PD in mA.
Output Power (W)	The power that is delivered to the PD in watts.
Status	The operational status of the port: • **Disabled**: No power is delivered. • **Delivering Power**. Power is being drawn by the PD. • **Requesting Power**: The port is requesting power. • **Fault**: A problem occurred with the power. • **Test**: The port is in test mode. • **Other Fault**: The port is idle because of an error condition. • **Searching**: The port is not in one of the other states in this list.
Fault Status	The error description when the PoE port is in a fault state: • **No Error**: The port is not in any error state and can provide power. • **MPS Absent**: The port detected the absence of the main power supply, preventing the port from providing power. • **Short**: The port detected a short circuit condition, preventing the port from providing power. • **Overload**: The PD that is connected to the port attempts to

Field	Description
	draw more power than allowed by the port's settings, preventing the port from providing power at all. • **Power Denied**: The port was denied power because of a shortage of power or because of an administrative condition. In this condition,

Power-cycle one or more PoE ports

You can power-cycle one or more PoE ports. This might be useful if PoE ports function not as expected.

To power-cycle one or more PoE ports:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > PoE > Advanced > PoE Port Configuration**.
   The PoE Port Configuration page displays.
3. Select one or more ports by taking one of the following actions:
   - To power-cycle a single port, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To power-cycle multiple ports with the same settings, select the check box associated with each port.
   - To power-cycle all ports with the same settings, select the check box in the heading row.
4. Click the **Power-Cycle Port(s)** button.
   The ports are power-cycled.

Manage the PoE usage threshold

If your model provides power supply unit (PSU) bays (also referred to as slots), you can install one or more auxiliary power supplies (APS) to expand PoE power or enable N+1 power redundancy (see Manage N+1 power redundancy on page 150). Depending on the number of PSU bays, you can also expand PoE power *and* enable N+1 power redundancy.
The PoE usage threshold applies to the switch, that is, to the internal power supply and any optional APSs that are installed.

To manage the PoE usage threshold for the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > PoE > Advanced > Power Configuration**.
   The Power Configuration page displays.
3. If a stack is configured, from the **UNIT ID** menu, select the ID for a stacked switch.
4. In the **System Usage Threshold** field, enter a number from 1 to 99 to set the threshold level at which a trap is sent if the consumed power exceeds the threshold power.
   The value that you can set in this field is the same value that you can set in the System Usage Threshold field on the PoE Configuration page (see Set the PoE system usage threshold and power management mode on page 141).
5. Click the **Apply** button.
   Your settings are saved.

   The following table describes the view-only fields in the Power Status section.

Table 51.Power status information

Field	Description
Total Power Available (W)	The maximum power in watts of all installed auxiliary power supplies that the switch can deliver to all ports
PoE Threshold Power (W)	If the consumed power is below the threshold power, the switch can power up another port. The consumed power can be between the nominal and threshold power. The threshold power is displayed in watts. **Note**: The threshold power value is determined by the value that you enter in the System Usage Threshold field.
Power Modules Slot	The PSU bay
Module Name	The APS model number
Status	The status of the APS: • **Not Present**: No APS is installed or detected • **Operational**: The APS is connected and works correctly • **Failed**: The switch cannot detect the status • **Not Powered**: The APS is connected but not powered on.
Power Module AC Input (V)	The APS AC input voltage

Manage N+1 power redundancy

If your model provides power supply unit (PSU) bays (also referred to as slots), you can install one or more auxiliary power supplies (APSs) to enable N+1 power redundancy or expand PoE power (see Manage the PoE usage threshold on page 148). Depending on the number of PSU bays, you can also do both.

When you enable N+1 power redundancy, the total usable power delivered by all APSs equals the power of all APSs minus the power of one APS. The power load is shared evenly across all APSs, which behave as if they were one large uninterruptible APS. If one APS fails, the two other APSs continue to provide power, allowing all connected powered devices (PDs) to remain powered up.

To manage N+1 power redundancy:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > PoE > Advanced > Power Configuration**.
   The Power Configuration page displays.
3. To enable or disable N+1 power redundancy, select one of the following radio buttons:
   - **Enable**: Enables N+1 power redundancy. The PDs might be powered off if the consumed power exceeds the available power.
   - **Disable**: Disables N+1 power redundancy. This is the default setting.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

   The following table describes the view-only fields in the Power Redundancy Configuration section.

Table 52.PoE information

Field	Description
N+1 Active	Displays whether N+1 is enabled (Yes or No)
Number of PSU	The number of connected and active APSs
Effective Number of PSU	The number of connected and active APSs while the N+1 mode is enabled

Display information about multiple power source management for PoE power

The default PoE power that a switch can deliver depends on the internal power supply. The *maximum* PoE power that a single switch can deliver depends on the number and types of optional auxiliary power supplies (APSs) that are installed in the switch.
For example, model M4350-48G4XF comes with one internal power supply unit (PSU) and provides two power supply bays that can accommodate different types of APSs (350W, 600W, 920W, or 2000W).
The multi-power source management (MPSM) table displays the available power for PoE operations for seven possible combinations out of these three power supplies (one internal power supply and two APS).

To display information about multiple power source management:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > PoE > Advanced > Power Configuration**.
   The Power Configuration page displays.

The following table describes the view-only fields in the Multiple Power Source Management section.

Table 53.Multiple Power Source Management information

Field	Description
Upper row	
Unit	The unit is always 1.
Slot	The slot is always 0.
MPSM	The MPSM index, which depends on the number of APSs that are installed in the switch: • 0: Abnormal condition • 1: Internal PSU only • 2: APS installed in PSU bay 1 • 3: Internal PSU and APS installed in PSU bay 1 • 4: APS installed in PSU bay 2 • 5: Internal PSU and APS installed in PSU bay 2 • 6: APS installed in PSU bay 1 and APS installed in PSU bay 2 • 7: Internal PSU, APS installed in PSU bay 1, and APS installed in PSU bay 2
MPSM Power Value	Total available power in watts based on the internal PSU and installed APSs
Lower row	
Unit	The unit is always 1.
Slot	The slot is always 0.
MPSM-0 (W)	Available power in watts (W), although the condition is not normal
MPSM-1 (W)	Available power in W from the internal PSU
MPSM-2 (W)	Available power in W from the APS installed in PSU bay 1
MPSM-3 (W)	Available power in W from the internal PSU and APS installed in PSU bay 1
MPSM-4 (W)	Available power in W from the APS installed in PSU bay 2
MPSM-5 (W)	Available power in W from the internal PSU and APS installed in PSU bay 2
MPSM-6 (W)	Available power in W from the APS installed in PSU bay 1 and APS installed in PSU bay 2

Field	Description
MPSM-7 (W)	Available power in W from the internal PSU, APS installed in PSU bay 1, and APS installed in PSU bay 2

Timer schedules

You can define multiple timer schedules (each with a unique name) that you can use for PoE power delivery to attached PDs and for access control lists:

- **Timers schedules for use with PoE**: After you create a timer schedule, you can associate it with one or more PoE ports (see Configure the PoE ports settings on page 142). You can also use a separate timer schedule for each PoE port.
  After you associate a timer schedule with a PoE port, the start date and time force the PoE port to *stop* delivering power and the stop date and time enable the PoE port to *start* delivering power.
- **Timers schedules for use with access control lists**: After you create a timer schedule, you can associate it with one or more access control lists (see Access control lists on page 685). You can also use a separate timer schedule for each access control list.

You can create absolute timer schedules, which apply to specific dates and times, and you can create recurring timer schedules. For each timer schedule, you can add multiple entries that apply to the selected timer schedule only.

Create a timer schedule

The maximum number of timer schedules that you can add is 100.

To create a timer schedule:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Timer Schedule > Basic > Global Configuration**.
   The Timer Schedule Name page displays.
3. In the **Timer Schedule Name** field, specify the name for a timer schedule.
4. Click the **Add** button.
   The timer schedule is added to the table on the Timer Schedule Name page and is assigned an ID.
5. To save the settings to the running configuration, click the **Save** icon.

Specify the settings for an absolute timer schedule

An absolute timer schedule applies to specific dates and times. The schedule is executed once only.

To specify the settings for a timer schedule that uses specific dates and times:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Timer Schedule > Advanced > Timer Schedule Configuration**.
   The Timer Schedule Configuration page displays.
3. In the Timer Schedule Selection section, make your selections from the following menus:
   a. **Timer Schedule Name**. Select the name of the timer schedule that you want to configure.
      You can select only names of schedules that you created (see Create a timer schedule on page 153).
   b. **Timer Schedule Type**. Select **Absolute**.
      The fields in the Timer Schedule Configuration section might adjust to let you configure a timer schedule for specific dates and times.
   c. **Timer Schedule Entry**. To add a new entry, select **new**.
      Selecting an existing entry lets you make changes to that entry.
4. In the Timer Schedule Configuration section, specify the times and dates:
   a. In the **Time Start** field, enter the time of day in the HH:MM format to specify when the timer schedule must start.
   b. In the **Time End** field, enter the time of day in the HH:MM format to specify when the timer schedule must stop.
   c. Next to the **Date Start** field, click the calendar icon and use the menus in the pop-up window to enter the date in the DD-Mon-YYY format to specify when the timer schedule must start.
   d. Next to the **Date End** field, click the calendar icon and use the menus in the pop-up window to enter the date in the DD-Mon-YYY format to specify when the timer schedule must stop.
5. Click the **Add** button.
   The entry for the timer schedule is added.
6. To save the settings to the running configuration, click the **Save** icon.

Specify the settings for a recurring timer schedule

A recurring schedule allows you to set up a single schedule that starts at a particular date and that recurs either with a specific end date or indefinitely.
For a single recurring timer schedule, you can add a daily, weekly, and monthly schedule configuration. That is, these schedule configurations are not mutually exclusive but complement each other.

To specify the settings for a timer schedule that uses a recurring pattern:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Timer Schedule > Advanced > Timer Schedule Configuration**.
   The Timer Schedule Configuration page displays.
3. In the Timer Schedule Selection section, make your selections from the following menus:
   a. **Timer Schedule Name**: Select the name of the timer schedule that you want to configure.
      You can select only names of schedules that you created (see Create a timer schedule on page 153).
   b. **Timer Schedule Type**: Select **Periodic**.
      The fields in the Timer Schedule Configuration section might adjust to let you configure a timer schedule with a recurrence pattern.
   c. **Timer Schedule Entry**: To add a new entry, select **new**.
      Selecting an existing entry lets you make changes to that entry.

4. In the Timer Schedule Configuration section, specify the recurrence pattern:
   a. In the **Time Start** field, enter the time of day in the HH:MM format to specify when the timer schedule must start.
   b. In the **Time End** field, enter the time of day in the HH:MM format to specify when the timer schedule must stop.
   c. Next to the **Date Start** field, click the calendar icon and use the menus in the pop-up window to enter the date in the DD-Mon-YYY format to specify when the timer schedule must start.
   d. Either select the **No End Date** radio button or select the **End Date** radio button, and next to the **End Date** field, click the calendar icon and use the menus in the pop-up window to enter the date in the DD-Mon-YYY format to specify when the timer schedule must stop.
   e. From the **Recurrence Pattern** menu, select the pattern:
      - **Daily**: The timer schedule works with daily recurrence. The fields adjust. Either select the **Every Weekday** radio button to let the schedule operate from Monday through Friday or select the **Every Day(s)** radio button and enter a number from 0 to 255 in the field.
        In the latter case, the schedule is triggered every specified number of days. If the number of days is not specified, or if you enter 0, then the schedule is triggered only once.
      - **Weekly**: The timer schedule works with weekly recurrence. The fields adjust. In the **Every Week(s)** field, enter a number from 0 to 255 to specify that the schedule must be triggered every specified number of weeks. If the number of weeks is not specified, or if you enter 0, then the schedule is triggered only once.
        Select a single **Week Day** check box, multiple check boxes, or all check boxes to specify the day or days of the week that the schedule must operate.
      - **Monthly**: The timer schedule works with monthly recurrence. The fields adjust. In the **Day** field, enter a number from 1 to 31 to specify the day of the month when the schedule must be triggered.
        In the **Every Month(s)** field, enter a number from 0 to 99 to specify that the schedule must be triggered every specified number of months. If the number of months is not specified, or if you enter 0, then the schedule is triggered only once.

5. Click the **Add** button.
   The entry for the timer schedule is added.

6. To save the settings to the running configuration, click the **Save** icon.

Change the settings for a recurring timer schedule entry

You can change the settings for an existing recurring timer schedule entry. (You cannot do this for an existing absolute timer schedule.)

To change the settings for an existing recurring timer schedule entry:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Timer Schedule > Advanced > Timer Schedule Configuration**.
   The Timer Schedule Configuration page displays.
3. From the **Timer Schedule Name** menu, select the schedule name.
4. From the **Timer Schedule Type** menu, select the schedule type.
5. From the **Timer Schedule Entry** menu, select the schedule entry.
6. Make the changes to the schedule entry.
   For more information, see Specify the settings for a recurring timer schedule on page 155.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Delete a timer schedule entry

You can delete a timer schedule entry that you no longer need.

To delete a timer schedule entry:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Timer Schedule > Advanced > Timer Schedule Configuration**.
   The Timer Schedule Configuration page displays.
3. From the **Timer Schedule Name** menu, select the schedule name.
4. From the **Timer Schedule Type** menu, select the schedule type.
5. From the **Timer Schedule Entry** menu, select the schedule entry.
6. Click the **Delete** button.
   The entry is deleted.

7. To save the settings to the running configuration, click the **Save** icon.

Delete a timer schedule

You can delete a timer schedule that you no longer need. All entries that are part of the timer schedule are also deleted.

To delete a timer schedule:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Timer Schedule > Basic > Global Configuration**.
   The Timer Schedule Name page displays.
3. Select the check box for the schedule that you want to delete.
4. Click the **Delete** button.
   The schedule is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Simple Network Management Protocol

You can configure SNMP settings for SNMPv1, SNMPv2, and SNMPv3. The switch supports the configuration of SNMP groups and users that can manage traps that the SNMP agent generates. The switch uses both standard public MIBs for standard functionality and private MIBs that support additional switch functionality.

Manage SNMPv1 and SNMPv2 communities

By default, no SNMP communities exist. The communities that you define can access the switch using SNMPv1 and SNMPv2. Only those communities with read/write level access can be used to change the configuration using SNMP.

Add an SNMPv1 and SNMPv2 community

You can add an SNMPv1 and SNMPv2 community, which allows both SNMPv1 and SNMPv2 access.

To add an SNMPv1 and SNMPv2 community:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **System > SNMP > SNMP V1/V2 > Community Configuration**.
   The Community Configuration page displays.
3. In the **Community Name** field, enter a name for a new community.
   The name can be up to 16 characters.
4. Specify the client IP address and client IP mask:
   - **Client Address**: Enter the IPv4 or IPv6 address of the client.
   - **Client IPv4 Mask/IPv6 Prefix Length**: Enter the IPv4 mask of an IPv4 client or the IPv6 prefix length of an IPv6 client.

   The client IP address and client IP mask or prefix length together denote a range of IP addresses from which SNMP clients can use the community to access the switch.
   For IPv4 clients, if either the client IP address or client IP mask is 0.0.0.0, access is allowed from any IP address. Otherwise, every client's IP address is ANDed with the mask, as is the client IP address. If the values are equal, access is allowed.
   For example, if the client IPv4 address and client IP mask are 192.168.1.0/255.255.255.0, any IPv4 client with an IP address in the range from 192.168.1.0 to 192.168.1.255 (inclusive) is allowed access. To allow access from only one IPv4 station, use a management station IP mask value of 255.255.255.255, and use that computer's IP address as the client address.
5. From the **Access Mode** menu, select the access level for this community, which is either **Read-Write** or **Read-Only**.
6. Click the **Add** button.
   The community is added.
7. To save the settings to the running configuration, click the **Save** icon.

Change an existing SNMPv1 and SNMPv2 community

You can change an existing SNMPv1 and SNMPv2 community.

To modify an existing SNMPv1 and SNMPv2 community:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > SNMP > SNMP V1/V2 > Community Configuration**.
   The Community Configuration page displays.
3. Select the check box next to the community.
4. Change the settings as needed.
   For more information about the settings, see Add an SNMPv1 and SNMPv2 community on page 158.

5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete an SNMPv1 and SNMPv2 community

You can delete an SNMPv1 and SNMPv2 community that you no longer need.

To delete an SNMPv1 and SNMPv2 community:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > SNMP > SNMP V1/V2 > Community Configuration**.
   The Community Configuration page displays.
3. Select the check box next to the community.
4. Click the **Delete** button.
   The community is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Manage the SNMPv1 and SNMPv2 trap settings

For each SNMP community, you can specify the source interface that must be used on the switch, the community name, the associated IP address, and other settings.

Add an SNMPv1 or SNMPv2 trap configuration for a host

You can add a trap configuration for a host, enabling the host to receive SNMPv1 or SNMPv2 traps.

To add an SNMPv1 or SNMPv2 trap configuration for a host:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > SNMP > SNMP V1/V2 > Trap Configuration**.
   The Trap Configuration page displays.

3. From the **Source Interface** menu, select the source interface that must be used for the SNMP trap receiver. By default, the following options display in the menu:
   - **None**: The primary IP address of the originating (outbound) interface is used as the source address.
   - **VLAN 1**: The primary IP address of VLAN 1 is used as the source address. This is the default selection.
   - **Service Port**: The management port IP address is used as the source address.

   Depending on the configuration of your switch, the following options can display:
   - **Another VLAN ID**: The primary IP address of a VLAN other than VLAN 1 is used as the source address.
   - **Routing interface**: The primary IP address of a routing interface is used as the source address.
   - **Routing VLAN**: The primary IP address of a VLAN routing interface is used as the source address.
   - **Routing loopback interface**: The primary IP address of a routing loopback interface is used as the source address.
   - **Different**: For some features, *Different* can display. This means that the source interface is configured separately.
4. In the **Community Name** field, specify the name of the SNMP community that includes the SNMP management host and the SNMP agent on the device.
   This name can be up to 16 characters and is case-sensitive.
5. From the **Version** menu, select the SNMP version to be used by the receiver:
   - **SNMPv1**: The switch uses SNMPv1 to send traps to the receiver. The default setting is SNMPv1.
   - **SNMPv2**: The switch uses SNMPv2 to send traps to the receiver.
6. From the **Protocol** menu, select the protocol to be used by the receiver:
   Select IPv4 if the receiver's address is an IPv4 address or IPv6 if the receiver's address is an IPv6 address. Or, you can select **DNS** and enter a domain name.
7. In the **Address** field, enter the IPv4 or IPv6 address or the domain name, depending on your selection from the **Protocol** menu.
8. From the **Status** menu, select to enable or disable the trap:
   - **Enable**: The trap is enabled.
   - **Disable**: The trap is configured but disabled.
9. Click the **Add** button.
   The trap configuration is added.
10. To save the settings to the running configuration, click the **Save** icon.

Change an SNMPv1 or SNMPv2 trap configuration for a host

You can change an existing SNMPv1 or SNMPv2 trap configuration for a host.

To change an SNMPv1 or and SNMPv2 trap configuration for a host:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > SNMP > SNMP V1/V2 > Trap Configuration**.
   The Trap Configuration page displays.
3. Select the check box next to the trap configuration.
4. Change the settings as needed.
   For more information about the settings, see Add an SNMPv1 or SNMPv2 trap configuration for a host on page 160.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete an SNMPv1 or SNMPv2 trap configuration for a host

You can delete an SNMPv1 or SNMPv2 trap configuration that you no longer need for a host.

To delete an SNMPv1 or SNMPv2 trap configuration for a host:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > SNMP > SNMP V1/V2 > Trap Configuration**.
   The Trap Configuration page displays.
3. Select the check box next to the trap configuration.
4. Click the **Delete** button.
   The trap configuration is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Configure SNMPv1 and SNMPv2 trap flags

You can enable or disable specific traps. When the condition that is identified by an active trap occurs on the switch, a trap message is sent to any enabled SNMP trap receivers, and a message is written to the trap log.

To configure the trap flags:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > SNMP > SNMP V1/V2 > Trap Flags**.
   The Trap Flags page displays.

3. For one or more trap flags that are displayed in the following table, select the **Disable** or **Enable** radio button.
   This selection enables or disables the activation of the trap.

Trap		Default	Behavior when enabled
Authentication		Enable	A trap is sent when an event involving authentication occurs, such as when a user attempts to access the switch main UI and does not provide a valid user name and password.
Link Up/Down		Enable	A trap is sent when the administrative or operational state of a physical or logical link changes.
Multiple Users		Enable	A trap is sent when the same user ID is logged into the switch more than once at the same time (either through Telnet or the serial port).
Spanning Tree		Enable	A trap is sent when a spanning tree change occurs.
ACL		Disable	A trap is sent when an ACL is triggered.
Captive Portal		Disable	A trap is sent when a user accesses a captive portal.
DVMRP		Disable	A trap is sent when a change in the DVMRP configuration occurs.
PIM		Disable	A trap is sent when a change in the PIM configuration occurs.
PoE (PoE models only)		Disable	A trap is sent when a change in the PoE configuration occurs.
OSPFv2 Traps			
Errors			
	Authentication-failure	Disable	A trap is sent when an OSPFv2 authentication failure occurs.
	Bad-packet	Disable	A trap is sent when a corrupt OSPFv2 packet is detected.
	Config-error	Disable	A trap is sent when an OSPFv2 configuration error occurs.
	Virt-authentication-failure	Disable	A trap is sent when a virtual OSPFv2 authentication failure occurs.

Trap		Default	Behavior when enabled
	Virt-bad-packet	Disable	A trap is sent when a corrupt virtual OSPFv2 packet is detected.
	Virt-config-error	Disable	A trap is sent when a virtual OSPFv2 configuration error occurs.
LSA			
	LSA-maxage	Disable	A trap is sent when the maximum age of an OSPFv2 link-state advertisement (LSA) is exceeded.
	LSA-originate	Disable	A trap is sent when an OSPFv2 LSA originates from the switch (or one of its routing interfaces), for example, when the topology changes.
Overflow			
	Lsdb-overflow	Disable	A trap is sent when The OSPFv2 link-state database (LSDB) is overflowing, that is, the LSDB can no longer be maintained because of resource constraints.
	Lsdb-approaching-overflow	Disable	A trap is sent when a critical threshold in the size of the OSPFv2 LSDB is exceeded.
Retransmit			
	Packets	Disable	A trap is sent when an OSPFv2 packet is retransmitted.
	Virt-packets	Disable	A trap is sent when a virtual OSPFv2 packet is retransmitted.
State-change			
	If-state-change	Disable	A trap is sent when the state of an OSPFv2 interface changes.
	Neighbor-state-change	Disable	A trap is sent when the state of an OSPFv2 neighbor changes.
	Virtif-state-change	Disable	A trap is sent when the state of a virtual OSPFv2 interface changes.
	Virtneighbor-state-change	Disable	A trap is sent when the state of a virtual OSPFv2 neighbor changes.
OSPFv3 Traps			

Trap		Default	Behavior when enabled
Errors			
	Bad-packet	Disable	A trap is sent when a corrupt OSPFv3 packet is detected.
	Config-error	Disable	A trap is sent when an OSPFv3 configuration error occurs.
	Virt-bad-packet	Disable	A trap is sent when a corrupt virtual OSPFv3 packet is detected.
	Virt-config-error	Disable	A trap is sent when a virtual OSPFv3 configuration error occurs.
LSA			
	LSA-maxage	Disable	A trap is sent when the maximum age of an OSPFv3 link-state advertisement (LSA) is exceeded.
	LSA-originate	Disable	A trap is sent when an OSPFv3 LSA originates from the switch (or one of its routing interfaces), for example, when the topology changes.
Overflow			
	Lsdb-overflow	Disable	A trap is sent when The OSPFv3 link-state database (LSDB) is overflowing, that is, the LSDB can no longer be maintained because of resource constraints.
	Lsdb-approaching-overflow	Disable	A trap is sent when a critical threshold in the size of the OSPFv3 LSDB is exceeded.
Retransmit			
	Packets	Disable	A trap is sent when an OSPFv3 packet is retransmitted.
	Virt-packets	Disable	A trap is sent when a virtual OSPFv3 packet is retransmitted.
State-change			
	If-state-change	Disable	A trap is sent when the state of an OSPFv3 interface changes.
	Neighbor-state-change	Disable	A trap is sent when the state of an OSPFv3 neighbor changes.

Trap		Default	Behavior when enabled
	Virtif-state-change	Disable	A trap is sent when the state of a virtual OSPFv3 interface changes.
	Virtneighbor-state-change	Disable	A trap is sent when the state of a virtual OSPFv3 neighbor changes.
Power Supply Module state		Enable	A trap is sent when the state of an auxiliary power supply (APS, also referred to as an external power supply unit [PSU]) changes. For example, a trap is sent if an APS is inserted, removed, changes to operational, or loses power, of if a warning is generated for the APS or the APS fails.
FAN Status		Enable	A trap is sent when the state of a fan module changes. For example, a trap is sent if a fan module is interted or removed, or if a warning is generated for the fan module or the fan module fails.
Temperature Status		Enable	A trap is sent when a critical temperature threshold is exceeded or if the temperature moves from its normal range.
VRRP		Enable	A trap is sent when one of the following conditions is detected: Invalid virutal router ID IP TTL error Invalid or unsupported version Checksum failure Invalid authentication type Authentication type mismatch from the VRRP packet when compared with the locally configured authentication type Authentication failure New VRRP master election
MAC Notification		Disable	A trap is sent when one of the following conditions is detected: New (MAC address) entry in the forwarding database (FDB) Removal of an existing entry from the FDB A change in an existing entry in the FDB

4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Display the supported MIBs

To display the MIBs supported by the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > SNMP > SNMP V1/V2 > Supported MIBs**.
   The Supported MIBs page displays.
   The Name field displays the RFC number, if applicable, and the name of the MIB.
   The Description field displays the RFC title or MIB description.

Manage SNMPv3 users

You can manage SNMPv3 user accounts and set up authentication and encryption.

Add an SNMPv3 user account

You can add an SMPv3 user account.

To add an SNMPv3 user account:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > SNMP > SNMP V3 > User Configuration**.
   The User Configuration page displays.
3. In the **User Name** field, type a name for the user account.
   The name can up be up to 30 characters.
4. From the **SNMP v3 Access Mode** menu, select **Read-Only** or **Read/Write** to specify the SNMPv3 access privileges for the user account.
   The SNMPv3 access privileges for the admin account are Read/Write.
5. From the **Authentication Protocol** menu, select **SHA512** for the authentication.
   SHA512 is the only possible choice from the menu.
6. In the **Authentication Key** field, type a password (key).
   The password can be up to 32 characters.

7. From the **Encryption Protocol** menu, select **None** or **AES128** for the encryption:
   - **None**: The information is not encrypted.
   - **AES128**: You must specify an encryption password for SNMPv3 access (see the following step).
8. If you select **AES128** from the **Encryption Protocol** menu, specify a password (key) in the **Encryption Key** field.
   The password can be up to 32 characters.
9. Click the **Add** button.
   The SNMPv3 user account is added.
10. To save the settings to the running configuration, click the **Save** icon.

Change an SNMPv3 user account

You can change an existing SMPv3 user account.

To change an existing SNMPv3 user account:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > SNMP > SNMP V3 > User Configuration**.
   The User Configuration page displays.
3. Select the check box next to the SNMPv3 user account.
4. Change the settings as needed.
   For more information about the settings, see Add an SNMPv3 user account on page 168.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete an SNMPv3 user account

You can delete an SMPv3 user account that you no longer need.

To delete an SNMPv3 user account:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **System > SNMP > SNMP V3 > User Configuration**.
   The User Configuration page displays.
3. Select the check box next to the SNMPv3 user account.
4. Click the **Delete** button.
   The SNMPv3 user account is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Link Layer Discovery Protocol

Link Layer Discovery Protocol (LLDP), which is the defined in IEEE 802.1AB, lets devices on a LAN advertise major capabilities and physical descriptions. You can view this information to identify the system topology and detect problematic configurations in the LAN.
LLDP is a one-way protocol without request and response sequences. Information is advertised by devices that are configured to transmit LLDP and is received and processed by devices that are configured to receive LLDP. You can enable and disable the transmit and receive functions separately per port. By default, both transmit and receive functions are disabled on all ports.

Configure the global LLDP settings

You can specify LLDP settings that are globally applied to the switch.

To configure the global LLDP settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > LLDP > Global Configuration**.
   The Global Configuration page displays.
3. In the **TLV Advertised Interval** field, enter the interval in seconds to transmit LLDP frames.
   The range is from 5 to 32768 secs. The default is 30 seconds.
4. In the **Hold Multiplier** field, enter the multiplier for the value that you enter in the **TLV Advertised Interval** field, which determines the time-to-live (TTL) for LLDP notifications. The range is from 2 to 10 secs. The default value is 4. As an example, if the value that you enter in the **TLV Advertised Interval** field is 30 and the value that you enter in the **Hold Multiplier** field is 4, the TTL for LLDP notifications is 120 seconds.
5. In the **Reinitialization Delay** field, enter the delay before reinitialization starts.
   The range is from 1 to 10 secs. The default is 2 seconds.

6. In the **NotificationInterval** field, enter the interval in seconds for the transmission of notifications.
   The range is from 5 to 3600 secs. The default is 5 seconds.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Configure LLDP interface settings

You can specify LLDP settings that are applied to one or more interfaces.

To configure LLDP interface settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > LLDP > Interface Configuration**.
   The Interface Configuration.page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

   The Link Status fields shows whether the interface is up or down.
5. From the **Transmit** menu, select if the interface can transmit LLDP notifications.
   The default is Enable.
6. From the **Receive** menu, select if the interface can receive LLDP notifications.
   The default is Enable.
7. From the **Notify** menu, select if the interface can support LLDP notifications.
   The default is Disable.

8. As an option, configure the following Type Length Value (TLV) information, all of which are enabled by default:
   - **Port Description**: Select if the interface can send a port description in an LLDP frame.
   - **System Name**: Select if the interface can send the system name in an LLDP frame.
   - **System Description**: Select if the interface can send the system description in an LLDP frame.
   - **System Capabilities**: Select if the interface can send the system capabilities in an LLDP frame.
9. From the **Management Information** menu, select if the management address is transmitted in LLDP frames.
10. Click the **Apply** button.
    Your settings are saved.
11. To save the settings to the running configuration, click the **Save** icon.

Display or clear LLDP statistics

To display or clear LLDP statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > LLDP > Statistics**.
   The Statistics page displays.
3. To refresh the page, click the **Refresh** button.
4. To clear all LLDP statistics, click the **Clear** button.
   All statistics are cleared.
5. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 54.LLDP statistics information

Field	Description
LLDP Statistics These statistics are for the switch.	
Last Update	The time when an entry was created, modified or deleted in the tables associated with the remote system.
Total Inserts	The number of times the complete set of information advertised by a particular MAC Service Access Point (MSAP) was inserted into tables associated with the remote systems.
Total Deletes	The number of times the complete set of information advertised by a particular MAC Service Access Point (MSAP) was deleted from tables associated with the remote systems.
Total Drops	The number of times the complete set of information advertised by a particular MAC Service Access Point (MSAP) could not be entered into tables associated with the remote systems because of insufficient resources.
Total Ageouts	The number of times the complete set of information advertised by a particular MAC Service Access Point (MSAP) was deleted from tables associated with the remote systems because the information timeliness interval has expired.
LLDP Statistics Therese statistics are for interfaces.	
Interface	The interface.
Transmit Total	The number of LLDP frames transmitted by the LLDP agent.
Receive Total	The number of valid LLDP frames received by this LLDP agent.
Discards	The number of LLDP TLVs discarded for any reason by the LLDP agent.
Errors	The number of invalid LLDP frames received by the LLDP agent.
Ageouts	The number of age-outs that occurred. An age-out is the number of times the complete set of information advertised by a particular MAC Service Access Point (MSAP) was deleted from tables associated with the remote entries because information timeliness interval expired.

Field	Description
TLV Discards	The number of LLDP TLVs discarded by the LLDP agent.
TLV Unknowns	The number of LLDP TLVs received that were not recognized by the LLDP agent.
TLV MED	The total number of LLDP-MED TLVs received.
TLV 802.1	The total number of LLDP TLVs received that are of type 802.1.
TLV 802.3	The total number of LLDP TLVs received that are of type 802.3.
TLV UPOE	The total number of LLDP TLVs received that are of type UPOE.
TLV NTGR	The total number of LLDP TLVs received that are of type NTGR.

Display LLDP local device information

You can display LLDP local device information, which is information that the switch itself, or an interface of the switch, advertises.

To display LLDP local device information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > LLDP > Local Device Information**.
   The Local Device Information page displays.
3. From the **Interface** menu, select an interface on which the transmission of LLDP frames is enabled.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 55.LLDP local device information

Field	Description
Chassis ID Subtype	The switch identifier is the MAC address of the switch (see the following field).
Chassis ID	The MAC address of the switch.
Port ID Subtype	The port identifier is the MAC address of the interface (see the following field).
Port ID	The MAC address of the interface on the switch.
System Name	The system name, if any, of the switch.
System Description	The description of the switch.
Port Description	The interface name, if any, of the interface on the switch.
System Capabilities Supported	The system capabilities of the switch.
System Capabilities Enabled	The system capabilities of the switch that are enabled.
Management Address Type	The type of the management IP address of the switch.
Management Address	The advertised management IP address of the switch.

Display LLDP remote device information

You can display LLDP remote device information, which is information that the device that is connected to an interface advertises.

To display LLDP remote device information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > LLDP > Remote Device Information**.
   The Remote Device Information page displays.
3. From the **Interface** menu, select an interface on which the reception of LLDP frames is enabled.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 56.LLDP remote device information

Field	Description
Remote ID	The ID of the remote device.
Switch ID	The MAC address of the remote device.
Switch ID Subtype	The switch identifier is the MAC address of the remote device (see the previous field).
Port ID	The identifier that is associated with the interface on the remote device.
Port ID Subtype	The type of the port identifier on the remote device.
System Name	The system name, if any, of the remote device.
System Description	The description of the remote device.
Port Description	The interface name, if any, of the interface on the remote device.
System Capabilities Supported	The system capabilities that are supported on the remote device.
System Capabilities Enabled	The system capabilities that are enabled on the remote device.
Time to Live	The time-to-live period in seconds of the received remote entry.
Management Address Type	The type of the management IP address of the remote device.
Management Address	The advertised management IP address of the remote device.

Display the LLDP remote device inventory

The LLDP inventory consists of the devices that LLDP detects.

To display the LLDP remote device inventory:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > LLDP > LLDP > Remote Device Inventory**.
   The Remote Device Inventory page displays.
3. To search for an interface, enter the interface number in the format 0/x, in which x is the interface number, in the **Search Interface** field and click the **Go** button.
   Information about the interface displays.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 57.LLDP remote device inventory information

Field	Description
Port	The interface on the switch on which LLDP information is received.
Remote Device ID	The remote device ID.
Management Address	The advertised management address of the remote device.
MAC Address	The MAC address of the remote device.
System Name	The system name of the remote device.
Remote Port ID	The interface number of the remote device.
OUI	The Type-Length-Value (TLV) organizationally unique identifier (OUI) of the remote device.
OUI Subtype	The subtype of the TLV OUI of the remote device.

Link Layer Discovery Protocol for Media Endpoint Devices

Link Layer Discovery Protocol for Media Endpoint Discovery (LLDP-MED) is an enhancement to LLDP with support for the following features:

- Auto-discovery of LAN policies (such as VLAN, Layer 2 priority, and DiffServ settings), which allows for plug-and-play networking.
- Device location discovery, which allows for the creation of location databases.
- Extended and automated power management of Power over Ethernet endpoints.
- Inventory management, which lets you track your network devices and determine their characteristics, such as manufacturer, software and hardware versions, and serial or asset numbers.

Configure the global LLDP-MED settings

You can specify LLDP-MED settings that are globally applied to the switch.

To configure the global LLDP-MED settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > LLDP > LLDP-MED > Global Configuration**.
   The Global Configuration page displays.

3. In the **Fast Start Repeat Count** field, enter the number of LLDP PDUs that are transmitted when the protocol is enabled.
   The range is from (1 to 10). Default value of fast repeat count is 3.
   The Device Class field displays the switch MED classification. The switch is a Network connectivity device.
   The following four types of MED devices exist, of which the first three are endpoints:
   - **Class I**: Generic devices, which include IP communication controllers.
   - **Class II**: Media devices, which include conference bridges.
   - **Class III**: Communication devices, which include IP telephones.
   - **Class IV**: Network connectivity devices, which include LAN switches, routers, IEEE 802.1 bridges, and IEEE 802.11 WiFi access points.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure LLDP-MED interface settings

You can specify LLDP-MED settings that are applied to one or more interfaces.

To configure LLDP-MED interface settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > LLDP > LLDP-MED > Interface Configuration**.
   The Interface Configuration page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.

4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

   The following two fields are view-only fields:
   - The Link Status fields shows whether the interface is up or down.
   - The Operational Status field shows whether LLDP-MED TLVs are transferred on this interface.
5. From the **MED Status** menu, select if LLDP-MED mode is enabled or disabled on the interface.
   The default is Enable.
6. From the **Notification Status** menu, select if the interface can support LLDP-MED notifications.
   The default is Disable.
7. Configure if the following transmit Type Length Value (TLV) information, all of which is enabled by default, is transmitted for LLDP-MED in LLDP PDU frames:
   - **MED Capabilities**. Select if the interface can send the MED capabilities in an LLDP frame.
   - **Network Policy**: Select if the interface can send the network policy in an LLDP frame.
   - **Location Identification**: Select if the interface can send location information in an LLDP frame.
   - **Extended MDI-PSE**: Select if the interface can send power sourcing equipment (PSE) information in an LLDP frame.
   - **Inventory Information**: Select if the interface can send inventory information in an LLDP frame.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Display LLDP-MED local device information

You can display LLDP-MED local device information, which is information that the switch itself, or an interface of the switch, advertises.

To display LLDP-MED local device information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > LLDP > LLDP-MED > Local Device Information**.
   The Local Device Information page displays.
3. From the **Interface** menu, select an interface on which the transmission of LLDP-MED frames is enabled.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 58.LLDP-MED local device information

Field	Description
Network Policy Information Displays if a network policy TLV is present in LLDP-MED frames that are transmitted.	
Media Application Type	The application type, which can be one of the following: ○ unknown ○ voice signaling ○ guest voice ○ guest voice signaling ○ soft phone voice ○ videoconferencing ○ streaming video ○ video signaling Each application type that is transmitted includes the VLAN ID, priority, DSCP, tagged bit status, and unknown bit status. An interface can receive one or many such application types. This information is displayed only if a network policy TLV is transmitted.
VLAN ID	The VLAN ID that is transmitted.
Priority	The priority that is transmitted.
DSCP	The DSCP that is transmitted.
Unknown Bit Status	The unknown bit status that is transmitted.
Tagged Bit Status	The tagged bit status that is transmitted.
Inventory Information Displays if an inventory TLV is present in LLDP-MED frames that are transmitted.	
Hardware Revision	The hardware version that is transmitted.
Firmware Revision	The firmware version that is transmitted.
Software Revision	The software version that is transmitted.
Serial Number	The serial number that is transmitted.
Manufacturer Name	The manufacturers name that is transmitted.
Model Name	The model name that is transmitted.
Asset ID	The asset ID that is transmitted.

Field	Description
Location Information Displays if a location TLV is present in LLDP-MED frames that are transmitted.	
Sub Type	The type of location information that is transmitted.
Location Information	The location information that is transmitted.
Extended PoE Displays if PoE device information is present in LLDP-MED frames that are transmitted.	
Device Type	The type of PoE device that is transmitted. The switch can supply PoE, so it is a power sourcing equipment (PSE).
Power Source	The type of power source that is transmitted.
Power Priority	The type of power priority that is transmitted.
Power Value	The power in watts that is transmitted, that is, the information is transmitted, not the actual power.

Display the LLDP-MED remote device information

You can display LLDP-MED remote device information, which is information that the device that is connected to an interface advertises.

To display the LLDP-MED remote device information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > LLDP > LLDP-MED > Remote Device Information**.
   The Remote Device Information page displays.
3. From the **Interface** menu, select an interface on which the reception of LLDP-MED frames is enabled.
   The Remote ID field displays the ID that is assigned to the remote device.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 59.LLDP-MED remote device information

Field	Description
Capability Information Displays the supported capabilities that are received in LLDP-MED frames.	
Supported Capabilities	The system capabilities that are supported on the remote device.
Enabled Capabilities	The system capabilities that are enabled on the remote device.
Device Class	The device class that is advertised by the remote device. The device class can be Generic, Media, Communication, or Network Connectivity.
Network Policies Information Displays the network policies information that is received in LLDP-MED frames.	
Media Application Type	The application type of the remote device, which can be one of the following: ○ unknown ○ voice signaling ○ guest voice ○ guest voice signaling ○ soft phone voice ○ videoconferencing ○ streaming video ○ video signaling Each application type that is received includes the VLAN ID, priority, DSCP, tagged bit status, and unknown bit status. An interface can receive one or many such application types. This information is displayed only if a network policy TLV is received.
VLAN ID	The VLAN ID of the remote device.
Priority	The priority of the remote device.
DSCP	The DSCP of the remote device.
Unknown Bit Status	The unknown bit status of the remote device.
Tagged Bit Status	The tagged bit status of the remote device.
Inventory Information Displays the inventory information that is received in LLDP-MED frames.	

Field	Description
Hardware Revision	The hardware version of the remote device.
Firmware Revision	The firmware version of the remote device.
Software Revision	The software version of the remote device.
Serial Number	The serial number of the remote device.
Manufacturer Name	The manufacturers name of the remote device.
Model Name	The model name of the remote device.
Asset ID	The asset ID of the remote device.
Location Information Displays the location information that is received in LLDP-MED frames.	
Sub Type	The type of location information of the remote device.
Location Information	The location information of the remote device.
Extended PoE Displays if extended PoE information is received in LLDP-MED frames.	
Device Type	The PoE device type of the remote device.
Extended PoE PSE: Displays if extended PoE PSE information is received in LLDP-MED frames.	
Device Type	The type of PoE device that remote device is.
Power Source	The type of power source of that the remote device.
Power Priority	The type of power priority on the remote device.
Power Value	The power in watts that the remote device transmits, that is, the information is transmitted, not the actual power.
Extended PoE PD: Displays if extended PoE PD information is received in LLDP-MED frames.	
Device Type	The type of PoE device that remote device is.
Power Source	The type of power source of that the remote device.
Power Priority	The type of power priority on the remote device.
Power Value	The power in watts that the remote device transmits, that is, the information is transmitted, not the actual power.

Display the LLDP-MED remote device inventory

The LLDP inventory consists of the devices that LLDP detects and that support MED.

To display the LLDP-MED remote device inventory:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > LLDP > LLDP-MED > Remote Device Inventory**.
   The Remote Device Inventory page displays.
3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 60.LLDP-MED remote device inventory information

Field	Description
Port	The interface on the switch on which LLDP-MED information is received
Management Address	The advertised management address of the remote device
MAC Address	The MAC address of the remote device
System Model	The system model of the remote device
Software Revision	The software version that is running on the remote device

Link dependency

Link dependency lets you enable or disable one or more interfaces based on the link state of one or more *other* interfaces. That is, if you enable link dependency on an interface, the link state of that interface depends on the link state of another interface. For example, if port 0/2 depends on port 0/4 and the switch detects a link loss on port 0/4, the switch can automatically bring down the link on port 0/2. When the link is restored on port 0/4, the switch can automatically restore the link on port 0/2.
The other way around is also possible: If port 0/2 depends on port 0/4 and the switch detects a link loss on port 0/4, the switch can automatically bring up the link on port 0/2. When the link is restored on port 0/4, the switch can automatically bring down the link on port 0/2.

Configure a link dependency group

A link dependency group includes a list of upstream interfaces, a list of downstream interfaces, and an action that is configured for the group. If all upstream interfaces in the list go down, all downstream interfaces in the list are either brought down or brought up, depending on the action that you configure for the link dependency group.
After you configure the link dependency group, configure the upstream and downstream interfaces for the group (see Configure or display upstream and downstream interfaces for a link dependency group on page 187).

To configure a link dependency group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Link Dependency > Link Dependency Group Configuration**.
   The Link Dependency Group Configuration page displays.
3. Select one or more check boxes for the group IDs.
   The switch supports 16 group IDs.

   NOTE: The group IDs are not associated with interface numbers.

4. From the **Link Action** menu, specify the following actions that must occur on the downstream interfaces in the group when all upstream interfaces in the group are down:
   - **Link Down**: When all upstream interfaces in the group are down, all downstream interfaces in the group are brought down too. When any of the upstream interfaces in the group are up, all downstream interfaces in the group are brought up too.
   - **Link Up**: When all upstream interfaces in the group are down, all downstream interfaces in the group are brought up. When any of the upstream interfaces in the group are up, all downstream interfaces in the group are brought down.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Configure or display upstream and downstream interfaces for a link dependency group

In a link dependency group, if you add downstream interfaces to the group, the interfaces are brought down until you add one or more upstream interfaces to the group. Then, the link state of the downstream interfaces is determined by the link state of the upstream interfaces and the action that you configure for the group. We recommend that you first configure the upstream interfaces and then configure the downstream interfaces.
In one link dependency group, an interface can either be a member of the upstream list or the downstream list, bot not both.

To configure or display upstream and downstream interfaces for a link dependency group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Link Dependency > Link Dependency Interface Configuration**.
   The Link Dependency Interface Configuration page displays.
3. In the Link Dependency Group ID section, from the **Group ID** menu, select the group ID for which you want to display or configure the settings.

   NOTE: The group IDs are not associated with interface numbers.

4. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
5. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

6. From the **Downstream Interface** menu, select if the interface is a member of the group's downstream list.
   - **False**: The interface is not a member of the downstream list for the group. This is the default setting.
   - **True**: The interface is a member of the downstream list for the group.
7. Click the **Apply** button.
   Your settings are saved.
8. From the **Upstream Interface** menu, select if the interface is a member of the group's upstream list.
   - **False**. The interface is not a member of the upstream list for the group. This is the default setting.
   - **True**: The interface is a member of the up stream list for the group.
9. Click the **Apply** button.
   Your settings are saved.
   The Link Status field shows whether the link for the corresponding interface or LAG is up or down.
10. To save the settings to the running configuration, click the **Save** icon.
11. To refresh the page, click the **Refresh** button.
    The following table describes the view-only fields in the Link Dependency Group Statistic section.

Table 61.Link dependency group statistic information

Field	Description
Group ID	The group ID. The range is from 1 to 16.
Link Action	The action to be performed on downstream interfaces when all the interfaces in the upstream list go down.
Group State	The current state of the group.
Group Transitions	The number of group transitions.
Last Transition Time	The time of the last group transition.

Clear all interfaces in a link dependency group

You can clear all interfaces in a link dependency group. After you do so, the upstream list and downstream list for the group do not include any interfaces.

To clear all interfaces in a link dependency group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Link Dependency > Link Dependency Interface Configuration**.
   The Link Dependency Interface Configuration page displays.
3. In the Link Dependency Group ID section, from the **Group ID** menu, select the group ID for which you want to display or configure the settings.
   The range for the group ID is from 1 to 16.

   NOTE: The group IDs are not associated with interface numbers.
4. Click the **Clear** button.
   The interfaces are removed from the group.
5. To save the settings to the running configuration, click the **Save** icon.

Industry Standard Discovery Protocol

Industry Standard Discovery Protocol (ISDP) is a protocol that enables devices to share information in a network.
You can configure the global and interface settings for ISDP.

Configure the global ISDP settings

You can configure the ISDP settings that apply globally to the switch.

To configure the global ISDP settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > ISDP > Basic > Global Configuration**.
   The Global Configuration page displays.
3. Select the Admin mode **Disable** or **Enable** radio button to specify is the ISDP service is disable or enabled.
   The default is Enabled.
4. In the **Timer** field, specify the period in seconds between the transmission of ISDP packets.
   The range is from 5 to 254 seconds. The default is 30 seconds.

5. In the **Hold Time** field, specify the hold time for ISDP packets that the switch transmits.
   The hold time specifies how long a receiving device must store information sent in the ISDP packet before discarding it. The range from 10 to 255 seconds. The default is 180 seconds.
6. Select the Version 2 Advertisements **Disable** or **Enable** radio button to specify if the switch can send ISDPv2 packets.
   The default is Enabled.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 62.ISDP global configuration information

Field	Description
Neighbors table last time changed	The time that the information in the neighbors table was changed.
Device ID	The device ID of this switch.
Device ID Format Capability	The device ID format capability.
Device ID Format	The device ID format.

Configure ISDP settings for an interface

You can configure if an interface can communicate IDSP packets.

To configure ISDP settings for an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > ISDP > Advanced > Interface Configuration**.
   The Interface Configuration page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.

4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Admin mode** menu, select **Enable** or **Disable** to specify if the interface can communicate ISDP information.
   The default is Enable.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Display or clear ISDP neighbor information

You can view or clear the information about ISDP neighbors.

To display or clear ISDP neighbor information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > ISDP > Advanced > Neighbor**.
   The Neighbor page displays.
3. To search for a neighbor, from the Search menu, select **Device ID** or **Interface**, and in the field, enter the device ID or interface number.
4. To refresh the page, click the Refresh button.
5. To clear the ISDP neighbor information, click the **Clear** button.
   All information on the page is cleared.
6. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 63.ISDP neighbor information

Field	Description
Device ID	The device ID of the ISDP neighbor.
Interface	The interface on which the neighbor is discovered.
Address	The IP address of the neighbor.
Capability	The capability of the neighbor, which can be one of the following: • Router • Trans Bridge • Source Route • Switch • Host • IGMP • Repeater
Platform	The model type of the neighbor.
Port ID	The port ID on the neighbor.
Hold Time	The hold time for ISDP packets that the neighbor transmits.
Advertisement Version	The ISDP version sending from the neighbor.
Entry Last Changed Time	The time since the last entry was changed.
Software Version	The software version on the neighbor.

Display or clear ISDP statistics

You can display or clear the ISDP statistics.

To display or clear ISDP statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > ISDP > Advanced > Statistics**.
   The Statistics page displays.
3. To refresh the page, click the **Refresh** button.
4. To clear the ISDP statistics, click the **Clear** button.
   The statistics are cleared.

5. To save the settings to the running configuration, click the **Save** icon.

   The following table describes the view-only fields on the page.

Table 64.ISDP statistics information

Field	Description
ISDP Packets Received	The ISDPv2 and ISDPv3 packets received on the switch.
ISDP Packets Transmitted	The ISDPv2 and ISDPv3 packets transmitted on the switch.
ISDPv1 Packets Received	The ISDPv1 packets received on the switch.
ISDPv1 Packets Transmitted	The ISDPv1 packets transmitted on the switch.
ISDPv2 Packets Received	The ISDPv2 packets received on the switch.
ISDPv2 Packets Transmitted	The ISDPv2 packets transmitted on the switch.
ISDP Bad Header	The ISDP packets with bad headers received on the switch.
ISDP Checksum Error	The ISDP packets with checksum errors received on the switch.
ISDP Transmission Failure	The number of ISDP packets that the switch failed to transmit.
ISDP Invalid Format	The number of ISDP packets with an invalid format that the switch received.
ISDP Table Full	The size of the ISDP table.
ISDP IP Address Table Full	The size of the ISDP IP address table.

Configure the RGB LED color settings

You can configure individual RGB LED colors for each physical port on the front panel of the switch. Each color indicates which AV profile that is assigned to that port. You can also view the configured RGB LED colors in the Device View page of the main UI.

The following switch models support RGB LED color configuration:

- M4350-24F4X
- M4350-8M2V
- M4350-24M4X4V
- M4350-40F4C
- M4350-16M4V
- M4350-16C

To configure the RGB LED color settings for one or more ports:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **System > Management > RGB LED Color Configuration**.
   The RGB LED Color Configuration page displays.
3. To enable Global Color Mode, select the **Enable** radio button.
   By default, this setting is disabled.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. In the **Color** field, enter a color code or use the color picker to select a color.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.
   The color displays on the RGB LED under the port on the front panel of the switch and in the Device View page of the main UI.

Configure Switching Information

This chapter covers the following topics:

- VLANs
- Auto-VoIP
- Internet Small Computer System Interface
- Spanning Tree Protocol
- Multicast forwarding database
- Internet Group Management Protocol snooping
- Multicast Listener Discovery snooping
- Multicast VLAN registration
- MAC address table
- Port settings
- Link aggregation groups
- MLAG overview
- 802.1AS timing and synchronization
- Multiple Registration Protocol and 802.1Qav
- Loop protection

VLANs

Adding virtual LAN (VLAN) support to a Layer 2 switch offers some of the benefits of both bridging and routing. Like a bridge, a VLAN switch forwards traffic based on the Layer 2 header, which is fast, and like a router, it partitions the network into logical segments, which provides better administration, security, and management of multicast traffic.

By default, all ports on the switch are in the same broadcast domain. VLANs electronically separate ports on the same switch into separate broadcast domains so that broadcast packets are not sent to all the ports on the switch. When you set up a VLAN, users can be grouped by function instead of physical location.

Each VLAN in a network is assigned a VLAN ID, which appears in the IEEE 802.1Q tag in the Layer 2 header of packets transmitted on a VLAN. An end station can omit the tag, or the VLAN portion of the tag, in which case the first switch port to receive the packet can either reject it or insert a tag using its default VLAN ID. A port can handle traffic for more than one VLAN, but it can support only one default VLAN ID.

You can define VLAN groups to be stored in the VLAN membership table. Each switch can support up to 1024 VLANs. VLAN 1 is created by default for use as the management VLAN, and all ports are are assigned as members by default.

Manage the VLAN configuration on the switch

You can add, change, and delete VLANs, or reset the entire VLAN configuration on the switch to

the default settings.

Add a VLAN

You can add multiple VLANs to customize the switch for your network.
An internal VLAN is reserved by a port-based routing interface and is invisible to the end user. After an internal VLAN is allocated by the port-based routing interface, the VLAN cannot be assigned to a routing VLAN interface.

To add a VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Basic > VLAN Configuration**.
   The VLAN Configuration page displays.
3. In the VLAN Configuration section, specify the settings for the new VLAN:
   - **VLAN ID**: Specify the identifier for the new VLAN. The range of the VLAN ID can be from 1 to 4093.
   - **VLAN Name**: Type a name for new VLAN. The name can be up to 32 characters, including blanks.
     By default, the name for VLAN ID 1 is Default.

   NOTE: When you add a VLAN manually (as in this procedure), the VLAN Type field always shows Static. A VLAN that is created by GVRP registration initially uses a type of dynamic but you can change it to static (see Change a VLAN on page 196). The type of the default VLAN (VLAN ID 1) is always Default.
4. Click the **Add** button.
   The VLAN is added.
5. To save the settings to the running configuration, click the **Save** icon.

Change a VLAN

You can change the name for a statically or dynamically added VLAN. For a dynamically added VLAN (for example, a VLAN that is created by GVRP registration), you change the VLAN type to static.

To change a VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Basic > VLAN Configuration**.
   The VLAN Configuration page displays.
3. In the VLAN Configuration section, select the check box for the VLAN ID.
4. To change the VLAN name, in the **VLAN Name** field, type a name for new VLAN.
   The name can be up to 32 characters, including blanks.
5. To change the VLAN type from dynamic to static, from the **Make Static** menu, select **Enable**.
   A VLAN that is created by GVRP registration initially uses a type of dynamic but you can change it to static. You cannot change the type for a VLAN that you added manually. The type of the default VLAN (VLAN ID 1) is always Default.
6. Click the **Apply** button.
   You changes are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Delete one or more VLANs

You can delete one or more VLANs that you no longer need. You cannot delete the default VLAN (VLAN 1).

To delete one or more VLANs:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Basic > VLAN Configuration**.
   The VLAN Configuration page displays.
3. In the VLAN Configuration section, select the check boxes for the VLAN IDs.
4. Click the **Delete** button.
   The VLANs are deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Reset the entire VLAN configuration to default setting

You can reset all VLAN configuration settings on the switch to factory default settings, with the exception of the default VLAN (VLAN 1). The factory default values are as follows:

- All ports are assigned to the default VLAN of 1.
- All ports are configured with a PVID of 1.
- All ports are configured to an “Acceptable Frame Types value of Admit All Frames.”
- All ports are configured with ingress filtering disabled.
- All ports are configured to transmit only untagged frames.
- GVRP is disabled on all ports and all dynamic entries are cleared.

To reset the entire VLAN configuration to default settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > VLAN > Advanced > VLAN Configuration**.
   The VLAN Configuration page displays.
3. Select the **Reset Configuration** check box.
   A confirmation windows displays.

   WARNING: If you select this button and confirm your selection, all VLAN configuration settings on the switch are reset to their factory default values.

4. Click the **OK** button.
   All VLANs, except for the default VLAN, are deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Change the internal VLAN allocation settings

In most situations, you do not need to change the internal VLAN allocation settings.

To change the internal VLAN allocation settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Basic > VLAN Configuration**.
   The VLAN Configuration page displays.

3. In the **Internal VLAN Allocation Base** field, specify the VLAN allocation base for the routing interface.
   The default base range of the internal VLAN is from 1 to 4093.
4. Select the Internal VLAN Allocation Policy **Ascending** or **Descending** radio button to specify the policy for the internal VLAN allocation:
   - **Ascending**: VLANs are allocated in ascending order from 2 to 4093.
   - **Descending**: VLANs are allocated in descending order from 4093 to 2.

   VLAN 1 is the default VLAN.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Auto-Trunk overview

Auto-trunk is a feature that lets the switch automatically enable Trunk mode on capable physical links and LAG interfaces between partner devices. A trunk can carry all active VLANs. By default, the Auto-Trunk feature is enabled on the switch.
If the switch automatically configures a port as a trunk (that is, an Auto-Trunk), all VLANs on the switch become part of the trunk, allowing automatic configuration of all VLANs on the switch and on the partner device with which the trunk is established.
Before the switch configures an Auto-Trunk, the switch first detects the physical links with the partner device that also supports the Auto-Trunk feature, and then automatically configures the ports that are connected and capable of forming a trunk at both ends.
A trunk carries multiple VLANs and accepts both tagged and untagged packets. Typically, a connection between the switch and a partner device such as a router, access point, or another switch functions as a trunk.
For the switch to form an Auto-Trunk with a partner device, the following are required:

- The Auto-Trunk feature must be supported and globally enabled on the switch and the partner device. (On the NETGEAR switch, the Auto-Trunk feature is enabled by default.)
- The interconnected ports on both the switch and the partner device must be enabled. (On the NETGEAR switch, all ports are enabled by default.)
- LLDP must be enabled on the interconnected ports on both the switch and the partner device. (On the NETGEAR switch, LLDP is enabled by default on all ports.)
- The interconnected ports on the switch and the partner device must be in the default switch port mode, which is the General mode. If the ports are in the Access mode or already in the Trunk mode, an Auto-Trunk cannot be formed on an Auto-LAG.

For an Auto-Trunk, the PVID is automatically set to the management VLAN. If you want to change the PVID for an Auto-Trunk, change the management VLAN.
The Auto-Trunk feature functions together with the Auto-LAG feature (see Auto-LAG overview

on page 298). After an Auto-LAG is formed, the switch automatically applies trunk mode (that is, an Auto-Trunk) to the LAG at both ends. In other words, after an Auto-LAG is formed, the mode for the ports that participate in an Auto-LAG is automatically changed from the default switch port mode to the trunk port mode, and the Auto-LAG then becomes an Auto-Trunk.
After a port or an Auto-LAG becomes an Auto-Trunk, all VLANs on the switch become part of the trunk, and all VLANs on the switch and the partner device can be configured automatically.

Enable or disable Auto-Trunks

By default, the Auto-Trunk feature is globally enabled but you can globally disable it.

To enable or disable Auto-Trunks:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > VLAN Trunking Configuration**.
   The VLAN Trunking Configuration page displays.
3. In the Global Auto-Trunk Mode section, select the Admin Mode **Enable** or **Disable** radio button.
   By default, the Auto-Trunk feature is globally enabled.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the switch port mode settings for interfaces

You can configure switch port mode settings on interfaces. The switch port mode defines the purpose of the port based on the type of device it connects to and constraints the VLAN configuration of the port accordingly. Assigning the appropriate switch port mode helps simplify VLAN configuration and minimize errors.

The switch supports the following types of switch ports:

- **Access**: This mode is for ports connected to end devices. Access ports participate in one VLAN only. They accept both tagged and untagged packets, but always transmit untagged packets.
- **General**: This mode enables a custom configuration of a port. You can configure the general port VLAN attributes, such as membership (see Configure membership interfaces for a VLAN on page 203) and PVID, tagging, ingress filter, and so on (see Change the port VLAN ID settings on page 205). By default, all ports are initially configured in the general mode.
- **Trunk**: This mode is for ports that are connected to other switches or to a router. A trunk port can participate in multiple VLANs and accept both tagged and untagged packets.

In the following procedure, ports are referred to as interfaces.

To configure the switch port mode settings for interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > VLAN > Advanced > VLAN Trunking Configuration**.
   The VLAN Trunking Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.

4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Switch port Mode** menu, select one of the following:
   - **Access**: Select this mode if the interface is connected to end devices. Access interfaces participate in one VLAN only. They accept both tagged and untagged packets, but always transmit untagged packets.
   - **General**: Select this mode so that you can configure the interface. You can configure the general port VLAN attributes, such as membership (see Configure membership interfaces for a VLAN on page 203) and PVID, tagging, ingress filter, and so on (see Change the port VLAN ID settings on page 205). This is the default selection.
   - **Trunk**: Select this mode if the interface is connected to another switch or to a router. A trunk interface can participate in multiple VLANs and accept both tagged and untagged packets.
6. Do one of the following, depending on your selection in the previous step:
   - **Access**: If you selected **Access** from the **Switch port Mode** menu, from the **Access VLAN ID** menu, select the access VLAN for the interface.
   - **Trunk**: If you selected **Trunk** from the **Switch port Mode** menu, do the following:
     a. From the **Native VLAN ID** menu, select the native VLAN for the interface.
     b. In the **Trunk Allowed VLANs** field, enter the VLANs of which the interface can be a member when the switch port mode for the interface Trunk By default, the field includes all VLANs, even if they are not yet created. VLAN IDs are in the range from 1 to 4093. Use a hyphen (-) to specify a VLAN range, or a comma (,) to separate VLAN IDs in a list. Spaces are not permitted. A zero value clears the allowed VLANs. If you enter All, all VLANs in the range from 1 to 4093 are included.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Definition
Native VLAN Tagging	Displays if VLAN tagging is enabled: • **Disable**: When VLAN tagging is enabled, if the trunk port receives untagged frames, it forwards them on the native VLAN with no VLAN tag. • **Enable**: When VLAN tagging is disabled, if the trunk port receives untagged frames, it includes the native VLAN ID in the VLAN tag when it forwards frames.
Auto-Trunk Operational Status	Displays if Auto-Trunk is configured on the interface: • **Disable**: Auto-Trunk is not configured on the interface. • **Enable**: Auto-Trunk is configured on the interface.
Auto-Trunk Member VLAN List	The VLANs that are members of the Auto-Trunk, if an Auto-Trunk is configured on the interface.

Configure membership interfaces for a VLAN

To configure membership interfaces for a VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > VLAN Membership**.
   The VLAN Membership page displays.
   The ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed. In addition, all LAGs are displayed, whether or not a stack is configured.
3. From the **VLAN ID** menu, select the VLAN ID.
4. To add or remove all ports and LAGs simultaneously (you can also select them individually), select one of the following options from the **Group Operation** menu:
   - **Untag All**: All ports and LAGs are added as untagged members of the VLAN.
   - **Tag All**: All ports and LAGs are added as untagged members of the VLAN.
   - **Remove** All: All ports and LAGs are excluded from the VLAN

5. In the Ports table (or if a stack is configured, in one of the Ports tables), click a port once, twice, or three times to configure one of the following modes:
   - **T (tagged) member**: The port is added as a tagged member of the VLAN.
   - **U (untagged) member**: The port is added as an untagged member of the VLAN.
   - **Excluded member**: By default, the selection is blank, which means that the port is excluded from the VLAN. A port that is excluded can still be dynamically added to the VLAN through GVRP.

   If a stack is configured, you can select ports in multiple tables.
6. In the LAG table, click a LAG once, twice, or three times to configure one of the following modes:
   - **T (tagged) member**: The LAG is added as a tagged member of the VLAN.
   - **U (untagged) member**: The LAG is added as an untagged member of the VLAN.
   - **Excluded member**: By default, the selection is blank, which means that the LAG is excluded from the VLAN. A LAG that is excluded can still be dynamically added to the VLAN through GVRP.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 65.Advanced VLAN Membership

Field	Definition
VLAN Name	The name for the VLAN.
VLAN Type	The type of the VLAN you selected: • **Default** (VLAN ID = 1): Always present. • **Static**: A VLAN that you added manually. • **Dynamic**: A VLAN that was created through GVRP registration and that you did not convert to static, and that GVRP can therefore remove.

View the VLAN status on the switch

You can view the status of all configured VLANs.

To view the VLAN status on the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Switching > VLAN > Advanced > VLAN Status**.
   The VLAN Status page displays.
3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 66. VLAN Status

Field	Definition
VLAN ID	The VLAN identifier (VID) of the VLAN. The range of the VLAN ID is 1 to 4093.
VLAN Name	The name of the VLAN. VLAN ID 1 is always named Default.
VLAN Type	The type of the VLAN you selected: ○ **Default** (VLAN ID = 1): Always present. ○ **Static**: A VLAN that you added manually. ○ **Dynamic**: A VLAN that was created through GVRP registration and that you did not convert to static, and that GVRP can therefore remove.
Routing Interface	The interface associated with the VLAN, if VLAN routing is configured for this VLAN.
Member Ports	The interfaces and LAGs that are members of the VLAN.

Change the port VLAN ID settings

By default, each interface is assigned a port VLAN ID (PVID) of 1 because it is associated with the default VLAN, VLAN ID 1.
If you want to change the PVID for an interface, the interface must be a member of at least one other VLAN in addition to the default VLAN.
In addition to the PVID, you can configure other PVID-related settings.

To configure the PVID and PVID-related settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > Port PVID Configuration**.
   The Port PVID Configuration page displays.

3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. In the **PVID** field, specify the VLAN ID to assign to untagged or priority-tagged frames received on this interface.
   The default is 1.
6. In the **VLAN Member** field, specify the VLAN ID or list of VLANs of a member interface. VLAN IDs range from 1 to 4093. The default is 1. Use a hyphen (-) to specify a range or a comma (,) to separate VLAN IDs in a list. Spaces and zeros are not permitted.
7. In the **VLAN Tag** field, specify the VLAN ID or list of VLANs of a tagged interface. VLAN IDs range from 1 to 4093. Use a hyphen (-) to specify a range or a comma (,) to separate VLAN IDs in a list. Spaces and zeros are not permitted. To reset the VLAN tag configuration to the defaults, use the **None** keyword. Port tagging for the VLAN can be set only if the interface is a member of this VLAN.
8. From the **Acceptable Frame Types** menu, select the types of frames that can be received on the interface:
   - **VLAN only**: Untagged frames and priority-tagged frames received on the interface are discarded. VLAN-tagged frames are forwarded in accordance to the 802.1Q VLAN specification.
   - **Admit All**: Untagged frames and priority-tagged frames received on the interface are accepted and assigned the value of the port VLAN ID for the interface. VLAN-tagged frames are forwarded in accordance to the 802.1Q VLAN specification.
   - **Admit Untagged Only**: Untagged frames received on the interface are accepted. VLAN-tagged frames are forwarded in accordance to the 802.1Q VLAN specification.

9. From the **Ingress Filtering** menu, select one of the following options:
   - **Enable**: The frame is discarded if the interface is not a member of the VLAN with which this frame is associated. In a tagged frame, the VLAN is identified by the VLAN ID in the tag. In an untagged frame, the VLAN is the port VLAN ID of the interface that receives this frame.
   - **Disable**: All frames are forwarded in accordance with the 802.1Q VLAN bridge specification. The default is Disabled.
10. In the **Port Priority** field, specify the default 802.1p priority assigned to untagged packets arriving at the port.
    You can enter a number from 0 to 7.
11. Click the **Apply** button.
    Your settings are saved.
12. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 67.PVID configuration information

Field	Description
Current Ingress Filtering	Indicates whether ingress filtering is enabled for the interface.
Untagged VLANs	The number of untagged VLANs that the interface is a member of.
Tagged VLANs	The number of tagged VLANs that the interface is a member of.
Forbidden VLANs	The number of forbidden VLANs that the interface is a member of.
Dynamic VLANs	The number of dynamically added VLANs that the interface is a member of.

Configure a MAC-based VLAN

A MAC-based VLAN allows incoming untagged packets to be assigned to a VLAN and classify traffic based on the source MAC address of the packet.
You define a MAC-to-VLAN mapping by configuring an entry in the MAC-to-VLAN table. An entry is defined by its source MAC address and a VLAN ID. MAC-to-VLAN configurations are shared across all interfaces (that is, there is a system-wide table with MAC-address-to-VLAN-ID mappings).
When untagged or priority-tagged packets arrive at the switch and entries exist in the MAC-to-VLAN table, the switch attempts to find the source MAC address of the packet: If the switch finds an entry, the corresponding VLAN ID is assigned to the packet. If the packet is already priority-tagged, it maintains this value; otherwise, the priority is set to zero. The assigned VLAN ID is

verified against the VLAN table: If the VLAN is valid, ingress processing on the packet continues; otherwise the packet is dropped. This implies that you can manually configure a MAC-address-to-VLAN-ID mapping.

Add a MAC-based VLAN configuration

You can add a MAC-address-to-VLAN-ID mapping, which we refer to as MAC-based VLAN configuration.

To add a MAC-based VLAN configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > MAC Based VLAN**.
   The MAC Based VLAN page displays.
3. In the **MAC Address** field, type a MAC address that must be bound to a VLAN ID.
4. In the **VLAN ID** field, specify the VLAN ID in the range of 1 to 4093.
5. Click the **Add** button.
   The MAC-based VLAN configuration is added.
6. To save the settings to the running configuration, click the **Save** icon.

Delete a MAC-based VLAN configuration

To delete a MAC-based VLAN configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > MAC Based VLAN**.
   The MAC Based VLAN page displays.
3. Select the check box for the MAC-based VLAN configuration
4. Click the **Delete** button.
   The MAC-based VLAN configuration is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Configure an OUI-based VLAN

You can use an OUI-based VLAN to automatically add a VLAN to a port, based on the MAC OUI

of the incoming packet.
You define an OUI-based VLAN by configuring a VLAN that corresponds to the MAC OUI so that the switch learns the MAC address in the L2-MAC table. If the OUI matches, the incoming port becomes a member of the OUI-based VLAN. The incoming port PVID is then set to the OUI-based VLAN. When the MAC entry expires, the incoming port is removed from the OUI-based VLAN.

Add an OUI-based VLAN configuration

You can add an OUI-based VLAN configuration.

To add an OUI-based VLAN configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > OUI Based VLAN**.
   The OUI Based VLAN Configuration page displays.
3. In the **OUI** field, type an OUI that must be bound to a VLAN ID.
4. In the **VLAN ID** field, specify the VLAN ID in the range of 1 to 128.
5. Click the **Add** button.
   The OUI-based VLAN configuration is added.
6. To save the settings to the running configuration, click the **Save** icon.

Delete an OUI-based VLAN configuration

To delete an OUI-based VLAN configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > OUI Based VLAN**.
   The OUI Based VLAN Configuration page displays.
3. Select the check box for the OUI-based VLAN configuration
4. Click the **Delete** button.
   The OUI-based VLAN configuration is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Configure a protocol-based VLAN group

You can use a protocol-based VLAN to define filtering criteria for untagged packets. By default, if you do not configure a port-based (IEEE 802.1Q) or protocol-based VLAN, untagged packets are assigned to VLAN 1. You can override this behavior by defining either port-based VLANs or protocol-based VLANs, or both. Tagged packets are always handled according to the IEEE 802.1Q standard, and are not included in protocol-based VLANs.
If you assign an interface to a protocol-based VLAN for a specific protocol, untagged frames received on the interface for that protocol are assigned the protocol-based VLAN ID. Untagged frames received on the interface for other protocols are assigned the Port VLAN ID, either the default PVID (1) or a PVID you specifically assigned to the interface (see Change the port VLAN ID settings on page 205).
You define a protocol-based VLAN by creating a group. Each group has a one-to-one relationship with a VLAN ID, can include one to three protocol definitions, and can include multiple interfaces. When you create a group, you specify a name. A group ID is assigned automatically.

Add a protocol-based VLAN group

You can add a protocol-based VLAN group.

To add a protocol-based VLAN group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > Protocol Based VLAN Group Configuration**.
   The Protocol Based VLAN Group Configuration page displays.
3. In the **Group ID** field, specify a numerical ID.
   You can enter an ID in the range from 1 to 128.
4. In the **Group Name** field, specify a name for the new group.
   You can enter up to 16 characters.

5. In the **Protocol** field, specify one or more of the following protocols to be associated with the group:
   - **IP**: IP is a network layer protocol that provides a connectionless service for the delivery of data.
   - **ARP**: Address Resolution Protocol (ARP) is a low-level protocol that dynamically maps network layer addresses to physical medium access control (MAC) addresses.
   - **IPX**: The internetwork packet exchange (IPX) is a connectionless datagram network-layer protocol that forwards data over a network.

   Separate protocols by a comma (,). For example, to specify all three protocols, enter the following: ip,arp,ipx
   You can also enter hexadecimal or decimal values in the range of 0x0600(1536) to 0xFFFF(65535).
6. In the **VLAN ID** field, specify the VLAN ID.
   The ID must be a number in the range of 1 to 4093. An interface in the group assign this VLAN ID to untagged packets that the interface receives for the protocols that you include in this group.
7. Click the **Add** button.
   The protocol-based VLAN group is added.
   The Ports field displays the interfaces ports that belong to the group.
8. To save the settings to the running configuration, click the **Save** icon.

Change a protocol-based VLAN group

You can change an existing protocol-based VLAN group.

To change a protocol-based VLAN group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > Protocol Based VLAN Group Configuration**.
   The Protocol Based VLAN Group Configuration page displays.
3. Select the check box for the group ID.
4. Change the settings are needed.
   For more information, see Add a protocol-based VLAN group on page 210.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete a protocol-based VLAN group

You can delete a protocol-based VLAN group that you no longer need.

To delete a protocol-based VLAN group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > Protocol Based VLAN Group Configuration**.
   The Protocol Based VLAN Group Configuration page displays.
3. Select the check box for the group ID.
4. Click the **Delete** button.
   The protocol-based VLAN group is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Configure membership interfaces for a protocol-based VLAN group

For a protocol, an interface can belong to one protocol-based VLAN group only. If you already added an interface to a group for IP, you cannot add the interface to another group that also includes IP, but you *can* add it to a new group for IPX.

To configure membership interfaces for a protocol-based VLAN group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > Protocol Based VLAN Group Membership**.
   The Protocol Based VLAN Group Membership page displays.
   The ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed. In addition, all LAGs are displayed, whether or not a stack is configured.
3. From the **Group ID** menu, select the protocol-based VLAN group ID.
   The Group Name field shows the name for the protocol-based VLAN that you selected.
4. To display the current members of the selected protocol-based VLAN group, click the **Current Members** button.
   A pop-up window displays the members.

5. To select physical interfaces for the VLAN, do one the following:
   - **Select all physical interfaces**: Do one of the following:
     - **No switch stack**: Click the **Ports** icon above the Ports table.
     - **Switch stack**: If a switch stack is configured, click the **Unit ID** icon above Ports table for the stacked switch. The ID of the unit is a number from 1 to 8, representing the switch in the stack.
   - **Select individual physical interfaces**: Do one of the following:
     - **No switch stack**: In the Ports table, click one or more ports individually.
     - **Switch stack**: If a switch stack is configured, in the Ports table for the stacked switch, click one or more ports individually. You can select ports in multiple Ports tables.
6. To select LAGs for the VLAN, do one the following:
   - **Select all LAGs**: Click the **LAG** icon above the LAG table.
   - **Select individual LAGs**: In the LAG table, click one or more LAGs individually.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Configure an IP subnet-based VLAN

You configure an IP subnet-to-VLAN mapping by specifying a source IP address, network mask, and the VLAN ID. The IP subnet-to-VLAN configurations are shared across all interfaces of the switch.

Add an IP subnet-based VLAN

You can add an IP subnet-based VLAN.

To add an IP subnet-based VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > IP Subnet Based VLAN**.
   The IP Subnet Based VLAN page displays.
3. In the **IP Address** field, specify an IP address.
   Enter the IP address in dotted-decimal notation.
4. In the **Subnet Mask** field, specify a subnet mask that is associated with the IP address.
   Enter the subnet mask in dotted-decimal notation.

5. In the **VLAN ID** field, specify the VLAN ID to which the IP configuration must be bound.
   The VLAN ID can be in the range from 1 to 4093.
6. Click the **Add** button.
   The IP subnet-based VLAN is added.
7. To save the settings to the running configuration, click the **Save** icon.

Delete an IP subnet-based VLAN

You can delete an IP subnet-based VLAN that you no longer need.

To delete an IP subnet-based VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > IP Subnet Based VLAN**.
   The IP Subnet Based VLAN page displays.
3. Select the check box for the IP subnet-based VLAN.
4. Click the **Delete** button.
   The IP subnet-based VLAN is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Configure a double VLAN

A double VLAN (DVLAN) lets you configure an EtherType VLAN tag inside a VLAN tag. The EtherType VLAN tag can be an 802.1Q tag, virtual metropolitan area network (vMAN) tag, or custom tag.

To configure the global EtherType for the DVLAN and configure the DVLAN on one or more interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > Port DVLAN Configuration**.
   The Port DVLAN Configuration page displays.

3. From the **Global EtherType** menu, select a tag type that determines the first 16 bits of the DVLAN tag:
   - **802.1Q Tag**: The dot1q tag that represents 0x8100.
   - **vMAN Tag**: The virtual metropolitan area network (vMAN) tag that represents 0x88A8.
   - **Custom Tag**: A custom EtherType tag in the range from 0 to 65535. With this selection the Custom Value field displays.
4. if you select **Custom Tag** from the **Global EtherType** menu, enter the custom EtherType tag in the **Custom Value** field.
   The tag can be in the range from 0 to 65535.
5. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
6. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
7. From the **Admin Mode** menu, select **Enable** or **Disable** to specify if the selected DVLAN tag is added to frames that are processed on the interface.
   The default is Disabled and the selected DVLAN tag is not added to frames.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Configure a voice VLAN

You can configure a voice VLAN for use with voice traffic. For example, enable the voice VLAN on an interface that is connected to IP phones. A voice VLAN can ensure that the sound quality of IP phone traffic remains good when data traffic on the same interface is high.

To configure a voice VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > Voice VLAN Configuration**.
   The Voice VLAN Configuration page displays.
3. Select the Admin Mode **Disable** or **Enable** radio button to disable or enable the global voice VLAN mode on the switch.
   The default is Disable. If you want to use a voice VLAN, you must enable the global voice VLAN mode.
4. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.
5. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
6. From the **Interface Mode** menu, select the voice VLAN mode:
   - **Disable**: This is the default value.
   - **None**: Allows an IP phone to use its own configuration to send untagged voice traffic.
   - **VLAN ID**: An IP phone must be configured to send tagged voice traffic.
   - **Dot1p**: Configure voice VLAN 802.1p priority tagging for voice traffic. If you select this mode, enter the dot1p value in the **Value** field.
   - **Untagged**: An IP phone must be configured to send untagged voice traffic.

7. From the **CoS Override Mode** menu, select one of the following:
   - **Enable**: The interface ignores (that is, overrides) the 802.1p priority value in the Ethernet frames it receives from connected devices. The DSCP field overrides the CoS value.
   - **Disable**: The interface trusts the priority value in the received frame. The default is Disable.
8. From the **Authentication Mode** menu, select one of the following:
   - **Enable**: Voice traffic is allowed on an unauthorized voice VLAN interface. The default is Enable.
   - **Disable**: Voice traffic is allowed only an authorized voice VLAN interface, for which dot1x must be enabled (see Manage port authentication on individual ports on page 606).
9. In the **DSCP Value** field, configure the DSCP value for the interface.
   The range is from 0 to 64. The default is 0.
10. Click the **Apply** button.
    Your settings are saved.
    The Operational State field displays the operational status of the voice VLAN on the interface.
11. To save the settings to the running configuration, click the **Save** icon.

Configure Generic Attribute Registration Protocol

Generic Attribute Registration Protocol (GARP) allows network devices to share information such as VLAN IDs and multicast group membership across a bridged LAN. That is, GARP participants can register and deregister attribute values within the LAN. When a GARP participant declares or withdraws an attribute, the attribute value is recorded for that attribute and for the interface from which the declaration or withdrawal was made.
The following applies to GARP:

- Registration occurs only on interfaces that receive a GARP PDU with a declaration or withdrawal.
- Deregistration occurs only if all GARP participants that are connected to the same LAN segment as the interface withdraw the declaration.

Configure GARP switch settings

You can globally enable GARP VLAN registration protocol (GVRP) and GARP multicast registration protocol (GMRP) on the switch.

To configure GARP switch settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > GARP Switch Configuration**.
   The GARP Switch Configuration page displays.
3. Select the GVRP Mode **Disable** or **Enable** radio button to disable or enable GARP VLAN registration protocol (GVRP) globally for the switch.
   If GVRP is enabled, the switch can share VLAN IDs with devices in the network. The default is Disable.
4. Select the GMRP Mode **Disable** or **Enable** radio button to disable or enable GARP multicast registration protocol (GMRP) globally for the switch.
   If GVMP is enabled, the switch can share multicast information with devices in the network. The default is Disable.
5. Click the **Apply** button.
   Your settings are saved.
   It can take up to 10 seconds for GARP configuration changes to take effect.
6. To save the settings to the running configuration, click the **Save** icon.

Configure GARP settings for one or more interfaces

You can configure GARP settings for individual interfaces. These settings take effect only if GVRP mode, GMRP mode, or both modes are enabled on the switch.

To configure GARP settings for one or more interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> VLAN > Advanced > GARP Port Configuration**.
   The GARP Port Configuration page display.

3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **GVRP Mode** menu, select **Enable** or **Disable** to specify the GARP VLAN registration protocol mode for the port.
   If you select **Disable**, GVRP is not active and the join time, leave time, and leave all time options are without any effect. The default is Disable.
6. From the **GMRP Mode** menu, select **Enable** or **Disable** to specify the GARP multicast registration protocol mode for the port.
   If you select **Disable**, GMRP is not active and the join time, leave time, and leave all time options are without any effect. The default is Disable.
7. In the **Join Timer** field, specify the time in centiseconds between the transmission of GARP PDUs registering membership for a VLAN or multicast group.
   Enter a number between 10 and 100 (0.1 to 1.0 seconds). The default is 20 centiseconds (0.2 seconds). An instance of this timer exists for each GARP participant for each port.
8. In the **Leave Timer** field, specify the time in centiseconds to wait after receiving an unregister request for a VLAN or multicast group before deleting the associated entry. This allows time for another station to assert registration for the same attribute to maintain uninterrupted service. Enter a number between 20 and 600 (0.2 to 6.0 seconds). The default is 60 centiseconds (0.6 seconds). An instance of this timer exists for each GARP participant for each port.

9. In the **Leave All Timer** field, specify how frequently (in centiseconds) LeaveAll PDUs are generated.
   A LeaveAll PDU indicates that all registrations will be deregistered soon. To maintain registration, participants must rejoin. The leave all period timer is set to a random value in the range of LeaveAllTime to 1.5 * LeaveAllTime. The timer is specified in centiseconds. Enter a number between 200 and 6000 (2 to 60 seconds). The default is 1000 centiseconds (10 seconds). An instance of this timer exists for each GARP participant for each port.
10. Click the **Apply** button.
    Your settings are saved.
    It can take up to 10 seconds for GARP configuration changes to take effect.
11. To save the settings to the running configuration, click the **Save** icon.

Auto-VoIP

Voice over Internet Protocol (VoIP) enables telephone calls over a data network. Because voice traffic is typically more time-sensitive than data traffic, Auto-VoIP provides classification for voice packets so that they can be prioritized above data packets for better quality of service (QoS). With Auto-VoIP, voice prioritization is based on call-control protocols such as SIP, H.323, and SCCP on interfaces or organizationally unique Identifier (OUI) bits.

Configure Auto-VoIP protocol-based settings

To prioritize time-sensitive voice traffic over data traffic, protocol-based Auto-VoIP determines if packets carry the following VoIP protocols:

- Session Initiation Protocol (SIP)
- H.323
- Signalling Connection Control Part (SCCP)

VoIP packets that come in on an interface on which Auto-VoIP is enabled are marked with the specified CoS traffic class value.

To configure protocol-based settings for one or more interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Auto-VoIP > Protocol-based>Port Settings**.
   The Port Settings page displays.

3. In the Protocol Based Global Settings section, specify the following global settings:
   a. From the **Prioritization Type** menu, select **Traffic Class** or **Remark**.
      This specifies the type of prioritization.
   b. From the **Class Value** menu, specify the CoS class value to be reassigned for packets that the voice VLAN receives.
      You can select a value in the range from 0 to 6.

   In the Protocol Based Global Settings section, specify the Auto VoIP settings for the interfaces, which is described in the following steps.

4. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.

5. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

6. From the **Auto VoIP Mode** menu, select to enable or disable the Auto VoIP mode for the interface.
   Auto-VoIP is disabled by default.

7. Click the **Apply** button.
   Your settings are saved.
   The Operational Status field show if the interface is up or down.

8. To save the settings to the running configuration, click the **Save** icon.

Configure the Auto-VoIP OUI-based properties

With organizationally unique Identifier (OUI)–based Auto-VoIP, voice prioritization is provided based on OUI bits.

To configure the Auto-VoIP OUI-based properties:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Auto-VoIP > OUI-based > Properties**.
   The Properties page displays.
3. In the **Auto-VoIP VLAN ID** field, enter the VoIP VLAN ID of the switch.
   No default VLAN exists for Auto-VoIP, you must create a VLAN for Auto-VoIP.
4. From the **OUI-based priority** menu, select the OUI-based priority of the switch.
   The default value is 7.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Configure the OUI-based interface settings

You can configure the OUI interface settings.

To configure the OUI-based interface settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Auto-VoIP > OUI-based> Port Settings**.
   The page Port Settings displays.

3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Auto VoIP Mode** menu, select to enable or disable the Auto VoIP mode for the interface.
   Auto-VoIP is disabled by default.
6. Click the **Apply** button.
   Your settings are saved.
   The Operational Status field show if the interface is up or down.
7. To save the settings to the running configuration, click the **Save** icon.

Manage the OUI table

Device hardware manufacturers can include an OUI in a network adapter to help identify a hardware device. The OUI is a unique 24-bit number assigned by the IEEE registration authority. The switch comes preconfigured with the following OUIs that identify the IP phone manufacturer:

- 00:01:E3: SIEMENS
- 00:03:6B: CISCO1
- 00:12:43: CISCO2
- 00:0F:E2: H3C
- 00:60:B9: NITSUKO
- 00:D0:1E: PINTEL
- 00:E0:75: VERILINK
- 00:E0:BB: 3COM

- 00:04:0D: AVAYA1
- 00:1B:4F: AVAYA2
- 00:04:13: SNOM

You can select an existing OUI or add a new OUI and description to identify the IP phones on the network.

Add an OUI prefix

You can add an OUI prefix to the OUI table.

To add an OUI prefix:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Auto-VoIP > OUI-based > OUI Table**.
   The OUI Table page displays.
3. In the **Telephony OUI(s)** field, specify the new VoIP OUI prefix in the format AA:BB:CC.
   The table can include up to 128 OUIs.
4. In the **Description** field, enter the description for the OUI.
   The maximum length is 32 characters.
5. Click the **Add** button.
   The telephony OUI entry is added.
6. To save the settings to the running configuration, click the **Save** icon.

Delete one or more OUI prefixes

You can delete OUI prefixes that you no longer need.

To delete one or more OUI prefixes:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Auto-VoIP > OUI-based > OUI Table**.
   The OUI Table page displays.
3. Select the check box next to each OUI prefix to be removed.
4. Click the **Delete** button.
   The telephony OUI entries are removed.

5. To save the settings to the running configuration, click the **Save** icon.

Display the Auto-VoIP status

To display the Auto-VoIP status:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Auto-VoIP > Auto-VoIP Status**.
   The Auto-VoIP Status page displays.
3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 68. Auto-VoIP status information

Field	Description
Auto-VoIP VLAN ID	The Auto-VoIP VLAN ID.
Maximum Number of Voice Channels Supported	The maximum number of voice channels supported.
Number of Voice Channels Detected	The number of VoIP channels prioritized successfully.

Internet Small Computer System Interface

The Internet Small Computer System Interface (iSCSI) feature can help you to track iSCSI traffic between iSCSI initiators and target systems, usually storage systems such as network-attached storage (NAS) and storage area network (SAN) systems.
The switch monitors or snoops traffic to detect packets used by iSCSI devices in establishing iSCSI sessions and connections. Data from these exchanges is used to create classification rules that assign the traffic between the devices to a configured traffic class. Packets in the flow are queued and scheduled for egress on the destination port based on these rules.
In a network that includes iSCSI initiators and targets, iSCSI can monitor iSCSI sessions or give iSCSI traffic preferential Quality of Service (QoS) treatment. Dynamically generated classifier rules are used to direct the iSCSI data traffic to queues that you can assign preference over other data traveling through the switch. This can prevent session interruptions during times of congestion that would otherwise cause iSCSI packets to be dropped. However, in systems where a large proportion of traffic is iSCSI, it might also interfere with other network control-plane traffic, such as ARP or LACP. That means that you must balance the preferential treatment of iSCSI traffic with the requirements of other critical data in the network.

You can view and manage iSCSI optimization by giving traffic between an iSCSI initiator and a target system a special QoS treatment.
In addition, enabling the *remark* function causes packets to be updated with IEEE 802.1 or IP-DSCP values. Remarking packets with priority data provides special QoS treatment as the packets travel through the network.

Enable iSCSI and configure the QoS settings for iSCSI traffic

By default, iSCSI is disabled and the switch does not optimize iSCSI traffic. You can enable iSCSI and configure the QoS settings for iSCSI traffic.

To enable iSCSI and configure the QoS settings for iSCSI traffic:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > iSCSI > Basic > Global Configuration**.
   The iSCSI Global Configuration page displays.
3. Select an iSCSI Status radio button:
   - **Enable**: iSCSI is enabled for the switch.
   - **Disable**: iSCSI is disabled for the switch. This is the default setting.
4. Select a QoS Profile radio button and the associated settings:
   - **VLAN Priority Tag**: iSCSI flows are assigned to the highest VLAN priority tag. From the **VLAN Priority Tag** menu, select a priority tag in the range from 0 (the lowest priority) to 7 (the highest priority). The default tag is 5.
   - **DSCP**: iSCSI flows are mapped to the highest DSCP value that is not used for switch management or a voice VLAN. From the **DSCP** menu, select a DSCP value in the range from 0 to 63. The higher the value, the higher the priority. The default value is 46.
5. Select a Remark radio button:
   - **Disable**: iSCSI frames that leave the switch are *not* marked with the configured VLAN priority tag or DSCP value.
   - **Enable**: iSCSI frames that leave the switch are marked (that is, *remarked*) with the configured VLAN priority tag or DSCP value. This is the default setting.
6. In the **iSCSI Aging Time** field, enter the number of minutes after which an inactive session is terminated.
   The range is from 1 to 43200 minutes. The default is 10 minutes,
7. Click the **Apply** button.
   Your settings are saved.

8. To save the settings to the running configuration, click the **Save** icon.

Add an iSCSI target

You can add an iSCSI target such as a NAS or SAN system by assigning a target port, a port IP address, or a combination of both, and giving the target a name.

To add an iSCSI target:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > iSCSI > Advanced > iSCSI Targets**.
   The iSCSI Targets Configuration page displays.
3. In the **TCP Port** field, enter a TCP port.
   This is a TCP port number that is used for an iSCSI target. You can add up to 16 TCP port numbers.
   By default, iSCSI ports 860 and 3260 are preconfigured.
4. In the **IP Address** field, enter the IPv4 address of the iSCSI target.
5. In the **Target Name** field, enter a name for the iSCSI target.
   The name can have a length of up to 233 characters.
6. Click the **Add** button.
   Your settings are saved. The iSCSI target is added.
7. To save the settings to the running configuration, click the **Save** icon.

View iSCSI sessions

You can view information about active iSCSI sessions.

To view iSCSI sessions:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > iSCSI > Advanced > Sessions**.
   The iSCSI Sessions page displays. If no sessions are present, the page displays *No iSCSI Sessions Are Available*.
3. To refresh the page, click the **Refresh** button.

The following table describes the view-only fields on the page.

Field	Description
Target Name	The target name.
Initiator Name	The initiator name.
Initiator Session ID (ISID)	The iSCSI session identifier. This is the ID that an initiator assigns to its session endpoint.

View detailed information about iSCSI sessions

You can view detailed information about active iSCSI sessions.

To view detailed information about iSCSI sessions:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > iSCSI > Advanced > Sessions Detailed**.
   The iSCSI Sessions Detailed page displays.
3. From the **Session Index** menu, select a session.
4. To refresh the page, click the **Refresh** button.

The following table describes the view-only fields on the page.

Field	Description
Target Name	The target name.
Initiator Name	The initiator name.
Up Time	The time in days, hours, minutes, and seconds since the current session started.
Time for Aging Out (in Seconds)	The time left in seconds before the session expires.
Initiator Session ID (ISID)	The iSCSI session identifier. This is the ID that an initiator assigns to its session endpoint.
Initiator IP Address	The IP address of the initiator.
Initiator TCP Port	The TCP port number that is used by the initiator. (More than one connection can exist between the initiator and the target.)
Target IP Address	The IP address of the target.
Target TCP Port	The TCP port number that is used by the target. (More than one connection can exist between the target and the initiator.)

Spanning Tree Protocol

The Spanning Tree Protocol (STP) provides a tree topology for any arrangement of network devices. STP also provides one path between end stations on a network, eliminating loops. STP (also referred to as “classic” STP) provides a single path between end stations, avoiding and eliminating loops. For information about configuring the global STP settings for the switch, see Configure the STP settings and display the STP status on page 230.

The switch support the following spanning tree versions:

- **CST**: Common STP. For information about configuring CST, see Configure the CST settings and display the CST status on page 232 and Configure the CST interface settings on page 235.
- **MSTP**: Multiple Spanning Tree Protocol (MSTP, also referred to as MST) supports multiple instances of spanning tree to efficiently channel VLAN traffic over different interfaces. For information about configuring MSTP, see Manage MST instances on page 240 and Configure and display the interface settings for an MST instance on page 242.
- **RSTP**: Rapid STP. Each instance of the spanning tree behaves in the manner specified in IEEE 802.1w, Rapid Spanning Tree (RSTP), with slight modifications in the working but not the end effect (the main effect is the rapid transitioning of the port to the forwarding state). The difference between the RSTP and the traditional STP (IEEE 802.1D) is the ability to configure

and recognize full-duplex connectivity and ports that are connected to end stations, resulting in rapid transitioning of the port to the forwarding state and the suppression of Topology Change Notification messages. These features are represented by the 'pointtopoint' and 'edgeport' parameters. MSTP is compatible with both RSTP and STP. It behaves in a way that is appropriate for STP and RSTP bridges. An MSTP bridge can be configured to behave entirely as an RSTP bridge or an STP bridge.

- **PVST**: Per-VLAN Spanning Tree (PVST) provides a unique instance of spanning tree for each VLAN. For information about configuring PVTP, see Configure the PVST/RPVST VLAN settings on page 246 and Configure the PVST and RPVST interface settings on page 248.
- **RPVST**: Rapid Per-VLAN Spanning Tree (RPVST) provides a unique instance of spanning tree for each VLAN with each instance behaving in the manner specified in IEEE 802.1w, Rapid Spanning Tree (RSTP). Compared to PVST, RPVST provides a faster spanning tree convergence in response to Layer 2 topology changes. For information about configuring RPVTP, see Configure the PVST/RPVST VLAN settings on page 246 and Configure the PVST and RPVST interface settings on page 248.

NOTE: For two bridges to be in the same region, the force version must be 802.1s and their configuration names, digest keys, and revision levels must match. For additional information about regions and their effect on network topology, refer to the IEEE 802.1Q standard.

Configure the STP settings and display the STP status

You can configure the Spanning Tree Protocol (STP) settings and display the STP status on the switch.

To configure the STP settings and display the STP status:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Basic > STP Configuration**.
   The STP Configuration page displays.
3. Select the Spanning Tree Admin Mode **Disable** or **Enable** radio button.
   This specifies whether spanning tree operation is enabled on the switch.

4. Select one of the following **Force Protocol Version** radio buttons to specify the STP version that the switch must use:
   - **IEEE 802.1d**: Traditional (classic) Spanning Tree Protocol.
   - **IEEE 802.1w**: Rapid Spanning Tree Protocol.
   - **IEEE 802.1s**: Multiple Spanning Tree Protocol.
   - **PVST**: Per-VLAN Spanning Tree Protocol.
   - **RPVST**: Rapid Per-VLAN Spanning Tree Protocol.
5. In the **Configuration Name** field, specify an ID or name that lets you identify the configuration.
   The name can be up to 32 characters. The default is the MAC address of the switch.
6. In the **Configuration Revision Level** field, specify an ID that lets you identify the configuration.
   For example, if you change the configuration, you can change the ID in the **Configuration Revision Level** field so you can keep track of the different configurations. You can enter an ID from 0 to 65535. The default is 0.
7. Select one of the following **Forward BPDU while STP Disabled** radio buttons:
   - **Disable**: Spanning tree BPDUs are not forwarded if spanning-tree is disabled on the switch. The default is Disable.
   - **Enable**: Spanning tree BPDUs are forwarded if spanning-tree is disabled on the switch.
8. Select one of the following **BPDU Guard** radio buttons:
   - **Disable**: The BPDU guard feature is disabled, which is the default setting.
   - **Enable**: The BPDU guard feature is enabled and enforces the STP domain borders to keep the active topology consistent and predictable. Switches behind edge ports on which STP BPDU guard is enabled do not affect the overall STP topology. At the reception of BPDUs, the port on which BPDU guard is enabled transitions into disabled state.
9. Select one of the following **BPDU Filter** radio buttons:
   - **Disable**: The BPDU Filter feature is disabled, which is the default setting.
   - **Enable**: The BPDU Filter feature is enabled and applies to all operational edge ports. An edge port in an operational state connects to hosts that typically drop BPDUs. If an operational edge port receives a BPDU, it loses its operational status. If a port on which BPDU filtering is enabled receives BPDUs, the port drops the BPDUs and remains operational.

10. Select one of the following **Fast Backbone** radio buttons:
    - **Disabled**: The Fast Backbone feature is disabled, which is the default setting.
    - **Enabled**: The Fast Backbone feature is enabled. If an indirect link fails, the switch can automatically select a new indirect link. With PVST, the switch does not ignore inferior BPDUs (which occurs with STP), but lets the BPDUs time-out on the port where they were received. Then, the switch sends root link queries on other non-designated ports. If the switch receives a positive response, it selects a new indirect link.
11. Select one of the following **Fast Uplink** radio buttons:
    - **Disabled**: The Fast Uplink feature is disabled, which is the default setting.
    - **Enabled**: The Fast Uplink feature is enabled. If the primary root port goes down, the recovery time in selecting a new root port is reduced.
12. If you enabled Fast Uplink in the previous step, in the **Max Update Rate** field, enter the maximum update rate in packets per second.
    You enter from 0 to 32000 packets per second. The default is 150.
13. Click the **Apply** button.
    Your settings are saved.
14. To save the settings to the running configuration, click the **Save** icon.
    The following table describes the view-only fields on the page.

Table 69.STP configuration and status information

Field	Description
STP Configuration	
Configuration Digest Key	The identifier for the STP configuration.
Configuration Format Selector	The version of the configuration format for the exchange of BPDUs.
STP Status	
MST ID	The multiple spanning tree (MST) or common spanning tree (CST) instance.
VID ID	The VLAN ID (VID).
FID ID	The filtering ID (FID).

Configure the CST settings and display the CST status

You can configure common spanning tree (CST) and display the CST status on the switch.

To configure the CST settings and display the CST status:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > CST Configuration**.
   The CST Configuration page displays.
3. In the **Bridge Priority** field, specify the bridge priority value for the common spanning tree (CST) and common and internal spanning tree (CIST).
   The range is from 0 to 61440. The bridge priority is a multiple of 4096. The default priority is 32768.
   When switches or bridges are running STP, each is assigned a priority. After exchanging BPDUs, the switch with the lowest priority value becomes the root bridge. If you specify a priority that is not a multiple of 4096, the priority is automatically set to the next lowest priority that is a multiple of 4096. For example, if you set the priority to a value between 0 and 4095, it is automatically set to 0.
4. In the **Bridge Max Age (secs)** field, specify the period in seconds that a bridge waits before implementing a topological change.
   The range is from 6 to 40 seconds, and the value must be less than or equal to the following: (2 * Bridge Forward Delay) – 1 and greater than or equal to 2 * (Bridge Hello Time +1).
   The default is 20 seconds.
5. In the **Bridge Hello Time (secs)** field, specify the period in seconds that a root bridge waits between configuration messages.
   The range is from 1 to 10. The default period is 2 seconds.
   The Bridge Hello Time is the interval at which the root bridge sends Bridge Protocol Data Units (BPDUs) to all other switches in the network.

   NOTE: This field is only editable on switches with firmware 14.0.9.x or above. On sites that do not meet these prerequisites, the Bridge Hello Time field is read-only.

6. In the **Bridge Forward Delay (secs)** field, specify the period in seconds that a bridge remains in a listening and learning state before forwarding packets.
   The period is from 4 to 30 seconds, and the value must be greater or equal to the following: (Bridge Max Age / 2) + 1.
   The default is 15 seconds.
7. In the **Spanning Tree Maximum Hops** field, specify the maximum number of bridge hops that the information for a particular CST instance can travel before being discarded.
   The range is from 6 to 40 hops. The default is 20 hops.

8. In the **Spanning Tree Tx Hold Count** field, specify the maximum number of BPDUs that the bridge can send within the hello time window.
   The range is from 1 to 10. The default is 6.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 70.CST status information

Field	Description
Bridge identifier	The bridge identifier for the CST, which consist of the bridge priority and the base MAC address of the bridge.
Time Since Topology Change	The time in seconds since the topology of the CST changed.
Topology Change Count	The number of times the topology changed for the CST.
Topology Change	This field shows if a topology change is in progress on any port assigned to the CST. (True or False.)
Designated Root	The bridge identifier of the root bridge, which consist of he bridge priority and the base MAC address of the bridge.
Root Path Cost	The path cost to the designated root for the CST.
Root Port	The port through which the designated root for the CST is accessed.
Max Age (secs)	The period in seconds that passes before a bridge port saves its configuration BPDU information.
Forward Delay (secs)	The period in seconds that the root port bridge delays the forwarding process. This period is a derived value.
Hold Time (secs)	The minimum period between the transmission of configuration BPDUs.
CST Regional Root	The priority and base MAC address of the CST regional root.
CST Path Cost	The path cost to the CST regional root.
Port Triggered TC	The port on which the topology change (TC) is triggered,

Configure the CST interface settings

You can configure a common spanning tree (CST) and internal spanning tree on a specific interface on the switch.
An interface can become diagnostically disabled (D-Disable) when a severe error condition occurs for DOT1S. The most common cause is when BPDU flooding occurs. The flooding criteria are such that DOT1S receives more than 15 BPDUs in a 3-second interval. (Other causes for a DOT1S D-Disable condition are very rare.)

To configure the CST interface settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > CST Port Configuration**.
   The CST Port Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. In the **Port Priority** field, specify the priority for the port in the CST.
   Specify the port priority in multiples of 16. For example, if you set the port priority to a value between 0 and 15, it is automatically set to 0. if you set the port priority to a value between 16 and (2*16-1), it is automatically set to 16, and so on. The default is 128.

6. From the **Admin Edge Port** menu, select **Enable** or **Disable** to specify if the port functions as an edge port in the CIST.
   The default is Disable.
7. In the **Port Path Cost** field, specify a path cost value for the port in the CIST
   The value can be in the range from 1 to 200000000. The default is 0.
8. From the **Auto Calculated Port Path Cost** menu, select **Enable** or **Disable** to specify if the external path cost is calculated automatically.
   If the value in the **Port Path Cost** field is zero and the path cost is calculated automatically, the value is based on the link speed of the port.
   The default is Disable.
9. In the **External Port Path Cost** field, specify an external path cost value for the port in the CIST.
   The value can be in the range from 1 to 200000000. The default is 0.
10. From the **Auto Calculated External Port Path Cost** menu, select **Enable** or **Disable** to specify if the external path cost is calculated automatically.
    If the value in the **External Port Path Cost** field is zero and the external path cost is calculated automatically, the value is based on the link speed of the port.
    The default is Disable.
11. From the **BPDU Filter** menu, select **Enable** or **Disable** to specify if BPDU traffic is filtered on the port when STP is enabled on the port.
    The default value is Disable.
12. From the **BPDU Forwarding** menu, select **Enable** or **Disable** to specify if BPDU traffic that arrives on the port can be flooded if STP is disabled on the port.
    The default value is Disable.
13. From the **Auto Edge** menu, select **Enable** or **Disable** to specify if the port can become an edge port if the port does not receive BPDUs for a period.
    The default is Enable.
14. From the **Root Guard** menu, select **Enable** or **Disable** to specify if the port can discard any superior information that it receives and, in this way, protect the root of the device from changing.
    If you enable this option, the port is placed in the discarding state and does not forward any packets. The default value is Disable.
15. From the **Loop Guard** menu, select **Enable** or **Disable** to specify if the port is protected from Layer 2 forwarding loops.
    If you enable this option, the port moves from the ‘listening/learning/forwarding’ state to the ‘STP loop inconsistent blocking state.’ The default is Disable

16. From the **TCN Guard** menu, select **Enable** or **Disable** to specify if the port is restricted from propagating any topology change information that it receives.
    The default is Disable.
17. From the **Port Mode** menu, select **Enable** or **Disable** to specify if STP is enabled or disabled on the port.
    The default value is Disable.
18. Click the **Apply** button.
    Your settings are saved.
19. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 71.CST port configuration information

Field	Description
Hello Timer	The setting of the Hello Timer for the CST. By default, the setting is 2.
BPDU Guard Effect	Indicates if an edge port that receives BPDU packets is enabled or disabled.
Port Forwarding State	Indicates the current STP state of the interface. If enabled, the port state determines what forwarding action is taken on traffic. Possible port states are as follows: • **Disabled**: STP is currently disabled on the interface. The interface forwards traffic while learning MAC addresses. • **Blocking**: The interface is currently blocked and cannot be used to forward traffic or learn MAC addresses. • **Listening**: The interface is currently in the listening mode. The interface cannot forward traffic nor can it learn MAC addresses. • **Learning**: The interface is currently in the learning mode. The interface cannot forward traffic. However, it can learn new MAC addresses. • **Forwarding**: The interface is currently in the forwarding mode. The interface can forward traffic and learn new MAC addresses. • **Manual forwarding**: The interface is currently in the manual forwarding mode. The interface can forward traffic and learn new MAC addresses.

Display the CST interface status

You can display the Common Spanning Tree (CST) status information for interfaces on the switch.

To display the CST interface status:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > CST Port Status**.
   The CST Port Status page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. To refresh the page, click the **Refresh** button.

The following table describes the view-only fields on the page.

Table 72.CST interface status information

Field	Description
Interface	The physical interface or LAG that is associated with the CST
Port ID	The port identifier for port in the CST, which is created from the port priority and the interface number
Port Forwarding State	Indicates the current STP state of the interface. If enabled, the port state determines what forwarding action is taken on traffic. Possible port states are as follows: • **Disabled**: STP is currently disabled on the interface. The interface forwards traffic while learning MAC addresses. • **Blocking**: The interface is currently blocked and cannot be used to forward traffic or learn MAC addresses. • **Listening**: The interface is currently in the listening mode. The interface cannot forward traffic nor can it learn MAC addresses. • **Learning**. The interface is currently in the learning mode. The interface cannot forward traffic. However, it can learn new MAC addresses. • **Forwarding**: The interface is currently in the forwarding mode. The interface can forward traffic and learn new MAC addresses. • **Manual forwarding**: The interface is currently in the manual forwarding mode. The interface can forward traffic and learn new MAC addresses
Port Role	Each MST bridge port that is enabled is assigned a port role for each spanning tree. The port role is one of the following: Root, Designated, Alternate, Backup, Master, or Disabled.
Designated Root	The root bridge for the CST, which is created from the bridge priority and the base MAC address of the bridge
Designated Cost	The path cost offered to the LAN by the designated port
Designated Bridge	The bridge identifier of the bridge with the designated port. This identifier is created from the bridge priority and the base MAC address of the bridge.
Designated Port	The port identifier on the designated bridge that offers the lowest cost to the LAN. This identifier is created from the port priority and the interface number of the port.

Field	Description
Topology Change Acknowledge	Indicates if the topology change acknowledgement flag is set for the next BPDU to be transmitted on the port (True or False)
Edge port	Indicates if the port is enabled as an edge port
Point-to-Point MAC	The point-to-point status, which indicates is the port's link is a point-to-point link (True) or not or (False)
CST Regional Root	The bridge identifier of the CST regional root. This identifier is created from the bridge priority and the base MAC address of the bridge.
CST Path Cost	The path cost to the CST regional root
Port Up Time Since Counters Last Cleared	The time the counters were last cleared
Loop Inconsistent State	Indicates if the interface is in a loop-inconsistent state (False or True)
Transitions Into Loop Inconsistent State	The number of times the interface transitioned into the loop-inconsistent state
Transitions Out Of Loop Inconsistent State	The number of times the interface transitioned out of the loop-inconsistent state

Manage MST instances

You can add, change, or delete multiple spanning tree (MST) instances on the switch. The MST instance consists of an ID, a priority value, and a VLAN ID.

Add an MST instance and display the MST status

You can add an MST instance and display the MST status.

To add an MST instance and display the MST status:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > MST Configuration**.
   The MST Configuration page display.

3. Configure the settings for the MST instance:
   - **MST ID**: In the **MST ID** field, specify the ID of the MST. The ID can be in the range from 1 to 4094.
   - **Priority**: In the **Priority** field, specify the bridge priority value for the MST instance. When switches or bridges are running STP, each is assigned a priority. After exchanging BPDUs, the switch with the lowest priority value becomes the root bridge. The bridge priority is a multiple of 4096. If you specify a priority that is not a multiple of 4096, the priority is automatically set to the next lowest priority that is a multiple of 4096. For example, if you set the priority to any value between 0 and 4095, the switch automatically sets the value to 0. The range is from 0 to 61440. The default is 32768.
   - **VLAN ID**: From the **VLAN ID** menu, select the VLAN that must be associated with the MST instance.
4. Click the **Add** button
   The MST is added.
5. To save the settings to the running configuration, click the **Save** icon.
6. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 73.MST configuration information

Field	Description
Bridge Identifier	The bridge identifier for the MST instance, which is created by using the bridge priority and the base MAC address of the bridge.
Last TCN	The time in seconds since the topology of the MST instance changed.
Topology Change Count	The number of times the topology changed for the MST instance.
Topology Change	This field shows if a topology change is in progress on an interface in the MST. (True or False.)
Designated Root	The bridge identifier of the root bridge, which is created by using the bridge priority and the base MAC address of the bridge
Root Path Cost	The path cost to the designated root for the MST instance.
Root Port	The port through which the designated root for the MST instance can be accessed.

Change an MST instance

You can change an existing MST instance.

To change an existing MST instance:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > MST Configuration**.
   The MST Configuration page display.
3. Select the check box that is associated with the MST instance.
4. Change the settings as needed.
   For more information, see Add an MST instance and display the MST status on page 240.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete an MST instance

You can delete an MST instance that you no longer need.

To delete an MST instance:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > MST Configuration**.
   The MST Configuration page display.
3. Select the check box that is associated with the MST instance.
4. Click the **Delete** button.
   The MST instance is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Configure and display the interface settings for an MST instance

You can configure and display the interface settings for a Multiple Spanning Tree (MST) instance. An interface can become diagnostically disabled (D-Disable) when a severe error condition occurs for DOT1S. The most common cause is when BPDU flooding occurs. The flooding criteria are such that DOT1S receives more than 15 BPDUs in a 3-second interval. (Other causes for a DOT1S D-Disable condition are very rare.)

To configure and display the port settings for an MST instance:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > MST Port Configuration**.
   The MST Port Configuration page displays.
3. From the **Select MST** menu, select an MST instance.
   For information about adding MST instances, see Manage MST instances on page 240.
4. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
5. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
6. Configure the MST values for the selected interface or interfaces:
   - **Port Priority**: In the **Port Priority** field, specify the priority for the interface in the MST instance. The port priority is set in multiples of 16. If you specify a value that is not a multiple of 16, the priority is automatically set to the next lowest priority that is a multiple of 16. For example, if you set a value between 0 and 15, the priority is set to 0. If you specify a number between 16 and 31, the priority is set to 16. Specify a value in the range from 0 to 240. The default is 128.
   - **Port Path Cost**: In the **Port Path Cost** field, specify the path cost in the range from 0 to 200000000. The default is 0.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

9. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 74.MST interface configuration information

Field	Description
Operational Port Path Cost	The path cost that the interface uses
Auto Calculated Port Path Cost	Indicates if the path cost is automatically calculated (Enable) or not (Disable). If enabled, the path cost is calculated based on the link speed of the port if the configured value for Port Path Cost is zero.
Port ID	The port identifier for the port in the MST instance, which is created by using the port priority and the interface number
Port Uptime Since Last Clear Counters	The time since the counters were cleared
Port Mode	Indicates if STP is enabled for the interface
Port Forwarding State	Indicates the current STP state of the interface. If enabled, the port state determines what forwarding action is taken on traffic. Possible port states are as follows: ○ **Disabled**: STP is currently disabled on the interface. The interface forwards traffic while learning MAC addresses. ○ **Blocking**: The interface is currently blocked and cannot be used to forward traffic or learn MAC addresses. ○ **Listening**: The interface is currently in the listening mode. The interface cannot forward traffic nor can it learn MAC addresses. ○ **Learning**: The interface is currently in the learning mode. The interface cannot forward traffic. However, it can learn new MAC addresses. ○ **Forwarding**: The interface is currently in the forwarding mode. The interface can forward traffic and learn new MAC addresses. ○ **Manual forwarding**: The interface is currently in the manual forwarding mode. The interface can forward traffic and learn new MAC addresses.
Port Role	Each MST bridge port that is enabled is assigned a port role for each spanning tree. The port role is one of the following: Root, Designated, Alternate, Backup, Master, or Disabled.
Designated Root	The bridge identifier of the root bridge, which is created by using the bridge priority and the base MAC address of the bridge

Field	Description
Designated Cost	The path cost offered to the LAN by the designated port
Designated Bridge	The bridge identifier of the bridge with the designated port. This identifier is created from the bridge priority and the base MAC address of the bridge
Designated Port	The port identifier on the designated bridge that offers the lowest cost to the LAN. This identifier is created from the port priority and the interface number of the port.
Loop Inconsistent State	Indicates if the interface is in a loop-inconsistent state (False or True)
Transitions Into Loop Inconsistent State	The number of times the interface transitioned into the loop-inconsistent state
Transitions Out Of Loop Inconsistent State	The number of times the interface transitioned out of the loop-inconsistent state

Display the STP interface statistics

You can display information about the number and type of bridge protocol data units (BPDUs) that are transmitted and received on each interface.

To display the STP interface statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > STP Statistics**.
   The STP Statistics page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 75.STP interface statistics information

Field	Description
STP BPDUs Received	The number of STP BPDUs received on the interface
STP BPDUs Transmitted	The number of STP BPDUs transmitted from the interface
RSTP BPDUs Received	The number of RSTP BPDUs received on the interface
RSTP BPDUs Transmitted	The number of RSTP BPDUs transmitted from the interface
MSTP BPDUs Received	The number of MSTP BPDUs received on the interface
MSTP BPDUs Transmitted	The number of MSTP BPDUs transmitted from the interface

Configure the PVST/RPVST VLAN settings

You can configure the Per-VLAN Spanning Tree Protocol (PVST) and Rapid PVST (RPVST) settings for a VLAN.
PVST and RPVST must be enabled. For more information, see Configure the STP settings and display the STP status on page 230.

To configure the PVST/RPVST VLAN settings for the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > PVST VLAN**.
   The PVST VLAN page displays.
3. In the **VLAN ID** field, specify a unique VLAN ID.
   You can specify only VLANs for which STP is enabled and for which PVST or RPVST is enabled.
4. From the **Root** menu, select **None**, **Primary**, or **Secondary**.:
   - **None**: The default is None. The switch does not function as a root bridge or a standby bridge for the VLAN.
   - **Primary**: The switch functions as a root bridge for the VLAN. The bridge priority is changed from the default value of 32768 to a lower value to ensure a root bridge configuration.
   - **Secondary**: The switch functions as a standby bridge for the VLAN. The bridge priority is changed from the default value of 32768 to a lower value to ensure a standby bridge configuration.

5. In the **Hello Time** field, enter the spanning tree hello time interval for the VLANs.
   The hello time is the interval between sending successive BPDUs. Enter a value in the range from 1 to 10 seconds. The default is 2 seconds.
6. In the **Forward Time** field, enter the spanning tree forwarding delay interval for the VLAN.
   This interval is the period during which the switch can detects state before transitioning an interface port to the forwarding state.The range is from 4 to 30 seconds. The default is 15 seconds.
7. in the **Max Age** field, enter the spanning tree maximum age time for the VLAN.
   This time is the period after which a bridge interface saves its configuration information. The range is from 6 to 40 seconds. The default is 20 seconds.
8. In the **Priority** field, enter the priority for the VLAN.
   The allowed values are between 0 and 61440. The valid values are listed in the following table.

0	4096	8192
12288	16384	20480
24576	28672	32768 (default)
36864	40960	45056
49152	53248	57344
61440	**[/normalized/concept/concept/task/taskbody/steps/step/info/table/tgroup/tbody/row/entry/p/clean {""}) (clean]**	**[/normalized/concept/concept/task/taskbody/steps/step/info/table/tgroup/tbody/row/entry/p/clean {""}) (clean]**

   The default value is 32768. If the value that you enter is not among the specified values, the value is automatically rounded off to the nearest valid value.
9. Click the **Add** button.
   PVST/RPVST VLAN configuration is added.
10. To save the settings to the running configuration, click the **Save** icon.

Change a PVST/RPVST VLAN configuration

You can change an existing PVST/RPVST VLAN configuration.

To change an existing PVST/RPVST VLAN configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > PVST VLAN**.
   The PVST VLAN page displays.
3. Select the check box that is associated with the PVST/RPVST VLAN configuration.
4. Change the settings as needed.
   For more information, see Configure the PVST/RPVST VLAN settings on page 246.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Remove a PVST/RPVST VLAN configuration

You can delete a PVST/RPVST VLAN configuration that you no longer need.

To delete a PVST/RPVST VLAN configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > PVST VLAN**.
   The PVST VLAN page displays.
3. Select the check box that is associated with the PVST/RPVST VLAN configuration.
4. Click the **Delete** button.
   The PVST/RPVST VLAN configuration is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the PVST and RPVST interface settings

You can configure the Per-VLAN Spanning Tree Protocol (PVST) and Rapid PVST (RPVST) settings for an interface.
PVST and RPVST must be enabled. For more information, see Configure the STP settings and display the STP status on page 230.

To configure the PVST and RPVST interface settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > PVST Interface**.
   The PVST Interface page displays.
3. From the **VLAN ID** menu, select a VLANs for which STP and PVST or RPVST is already enabled, or select **Other** to configure a by entering and a new VLAN ID.
4. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
5. In the **Priority** field, enter the priority for the interface.
   Set this value to a lower number to let the interface forward frames. The priority s used when the interface is configured as a point-to-point link. The value can be from 0 to 240. The value must be a multiple of 16.

0	16	32
48	64	80
96	112	128 (default)
144	160	176
192	208	224
240	–	–

   The default value is 128. If the value that you enter is not among the specified values, the value is automatically rounded off to the nearest valid value.
6. In the **Cost** field, enter the path cost from the interface to the root bridge.
   By default, the cost is not manually configured. You can enter a value from 1 and 200,000,000. Enter 0 to disable the setting. If the per-VLAN cost is not configured, the path cost value is set based on the link speed.

7. From the **Auto Calculated Cost** menu, select **Enable** or **Disable** to specify if the cost is automatically calculated.
   The default is Enable.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 76.PVST/RPVST interface configuration information

Field	Description
Role	The role of the interface, which can be Disabled, Root, Designated, Alternate, Backup, or Master **Note**: For an interface that is not member of the selected VLAN, the field is blank.
Status	The status of the interface, which can be Discarding, Learning, Forwarding, or Disabled **Note**: For an interface that is not member of the selected VLAN, the field is blank.

Display the PVST statistics

You can display the Per-VLAN Spanning Tree Protocol (PVST) statistics.
PVST must be enabled. For more information, see Configure the STP settings and display the STP status on page 230.

To display the PVST statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > STP > Advanced > PVST Statistics**.
   The PVST Statistics page displays.

3. To refresh the page, click the **Refresh** button.
   If PVST is enabled, the following table describes the view-only fields on the page.

Table 77.PVST statistics information

Field	Description
Fast Backbone	
Transitions via Fast Backbone	The number of fast backbone transitions
Inferior BPDUs received	The number of received inferior Bridge Protocol Data Units (BPDUs)
RLQ Request PDUs Received	The number of received Root Link Query (RLQ) request Protocol Data Units (PDUs)
RLQ Response PDUs Received	The number of received RLQ response PDUs
RLQ Request PDUs Sent	The number of transmitted RLQ request PDUs
RLQ Response PDUs Sent	The number of transmitted RLQ response PDUs
Fast Uplink	
Fast Uplink Transitions	The number of fast uplink transitions
Proxy Multicast Addresses Transmitted	The number of transmitted proxy multicast addresses
Name	The name of the VLAN on which PVST is enabled
Interface List	The interfaces that are transmitting traffic for the VLAN, either as tagged or as untagged frames

Multicast forwarding database

Multicast IP traffic is traffic that is destined to a host group. Host groups are identified by class D IP addresses, which range from 224.0.0.0 to 239.255.255.255.

Display the entries in the multicast forwarding database

The multicast forwarding database (MFDB) holds the port membership information for all active multicast address entries. The key for an entry consists of a combination of a VLAN ID and a MAC address. Entries can contain data for more than one protocol.

To display the entries in the multicast forwarding database:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Multicast > MFDB > MFDB Table**.
   The MFDB Table page displays.
3. To search for an entry, in the **Search by MAC Address**, enter a MAC address, and click the **Go** button.
   If the address exists, that entry is displayed. An exact match is required.
   The following table describes the view-only fields on the page.

Table 78.MFDB information

Field	Description
MAC Address	The multicast MAC address for which you requested data.
VLAN ID	The VLAN ID to which the multicast MAC address is related.
Component	The component that is responsible for this entry in the MFDB. The component can be IGMP snooping, GMRP, Static Filtering, or MLD snooping.
Type	The type of the entry. Static entries are those that you configure. Dynamic entries are added to the table as a result of a learning process or protocol.
Description	The description of this multicast table entry. The description can be Management Configured, Network Configured, or Network Assisted.
Forwarding Interfaces	The forwarding list that is derived from combining all the forwarding interfaces and removing the interfaces that are listed as the static filtering interfaces.

Remove the IGMP snooping entries from the multicast forwarding database

You can remove the IGMP snooping entries from the multicast forwarding database.

To remove the IGMP snooping entries from the multicast forwarding database:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Switching > Multicast > MFDB > MFDB Table**.
   The MFDB Table page displays.
3. Click the **Clear** button.
   Your settings are saved.
4. To save the settings to the running configuration, click the **Save** icon.

Remove all known multicast MAC entries from the multicast forwarding database

You can remove all known multicast MAC entries from the multicast forwarding database.

To remove all known multicast MAC entries from the multicast forwarding database:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Multicast > MFDB > MFDB Table**.
   The MFDB Table page displays.
3. Select the **Reset Known Multicast MAC Entries** check box.
   A confirmation pop-up window displays.
4. Select the **OK** button.
   The pop-up window closes.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Display the multicast forwarding database statistics

You can display the multicast forwarding database statistics

To display the multicast forwarding database statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Multicast > MFDB > MFDB Statistics**.
   The MFDB Statistics table displays.

3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 79.MFDB statistics information

Field	Description
Max MFDB Table Entries	The maximum number of entries that the multicast forwarding database table can hold.
Most MFDB Entries Since Last Reset	The largest number of entries that were present in the multicast forwarding database table since the last reset. This value is also known as the MFDB high-water mark.
Current Entries	The current number of entries in the multicast forwarding database table.

Internet Group Management Protocol snooping

Internet Group Management Protocol (IGMP) snooping allows a switch to forward multicast traffic intelligently. Multicast IP traffic is destined to a host group. Host groups are identified by class D IP addresses, which range from 224.0.0.0 to 239.255.255.255. Based on the IGMP query and report messages, the switch forwards traffic only to the ports that request the multicast traffic. This prevents the switch from broadcasting the traffic to all ports and possibly affecting network performance.

A traditional Ethernet network can be separated into different network segments to prevent placing too many devices onto the same shared media. Bridges and switches connect these segments. When a packet with a broadcast or multicast destination address is received, the switch forwards a copy into each of the remaining network segments in accordance with the IEEE MAC Bridge standard. Eventually, the packet is made accessible to all nodes connected to the network.

This approach works well for broadcast packets that are intended to be seen or processed by all connected nodes. In the case of multicast packets, however, this approach could lead to less efficient use of network bandwidth, particularly when the packet is intended for only a small number of nodes. Packets are flooded into network segments where no node is receptive to the packet. While nodes rarely incur any processing overhead to filter packets addressed to unrequested group addresses, they cannot transmit new packets onto the shared media during the period that the multicast packet is flooded. The problem of wasting bandwidth is even worse when the LAN segment is not shared, for example in full-duplex links.

Allowing switches to snoop IGMP packets is a creative effort to solve this problem. The switch uses the information in the IGMP packets as they are being forwarded throughout the network to determine which segments receive packets directed to the group address.

Configure IGMP snooping automatically with IGMP Plus mode

IGMP Plus mode lets you automatically configure IGMP snooping, which is used to build forwarding lists for IPv4 multicast traffic.
You can also configure IGMP snooping manually (see Configure IGMP snooping manually on page 256).

To configure IGMP snooping automatically:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > IGMP Snooping > Configuration**.
   The Configuration page displays.
3. Select the IGMP Plus Mode **Enable** or **Disable** radio button.
   If enabled, the following IGMP snooping modes are automatically enabled:
   - Admin mode
   - Proxy Querier mode
   - Report Flood Mode
   - Exclude Mrouter Interface Mode
   - Fast Leave Auto-assignment Mode

   The default is Enable.
   If disabled, these IGMP snooping modes are automatically disabled.

   NOTE: For information about other settings on the page, see Configure IGMP snooping manually on page 256.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

6. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 80.IGMP snooping configuration information

Field	Description
Multicast Control Frame Count	The number of multicast control frames that are processed by the switch.
Interfaces Enabled for IGMP Snooping	The interfaces on which IGMP snooping is enabled.
Operational Mode	Indicates if IGMP snooping is globally enabled or disabled on the switch.
VLAN IDs Enabled For IGMP Snooping	The VLANs on which IGMP snooping is enabled.

Configure IGMP snooping manually

You can manually configure the settings for IGMP snooping, which is used to build forwarding lists for IPv4 multicast traffic.
You can also configure IGMP snooping automatically (see Configure IGMP snooping automatically with IGMP Plus mode on page 255).

To configure the settings for IGMP snooping manually:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > IGMP Snooping > Configuration**.
   The Configuration page displays.
3. Select the Admin Mode **Enable** or **Disable** radio button.
   This selection enables or disables the administrative mode for IGMP snooping for the switch. The default is Disable.
4. Select the Validate IGMP IP header **Enable** or **Disable** radio button.
   This selection enables or disables header validation for all IGMP versions on the switch. If enabled, a packet IGMP IP header validates the Router Alert option, ToS and TTL. The default is Enable.

5. Select the Proxy Querier Mode **Enable** or **Disable** radio button.
   This selection enables or disables the IGMP proxy querier for the switch. If disabled, the IGMP proxy query with source IP address 0.0.0.0 is not sent in response to an IGMP leave packet. The default is Enable.
6. Select the Report Flood Mode **Enable** or **Disable** radio button.
   This selection enables or disables the report flooding mode on the switch. If enabled, IGMP Join/Leave PDUs that the switch receives from a host on a downstream port are forwarded to all other ports in the associated VLAN. The default is Enable.
7. Select the Exclude Mrouter Interface Mode **Enable** or **Disable** radio button.
   This selection specifies the type of information that is forwarded to the upstream multicast router interface.
   If enabled, the switch forwards IGMP Join/Leave PDUs that it receives on a downstream port to an upstream mrouter interface. In addition, the switch forwards a multicast data stream to an upstream mrouter interface only if that port already received an IGMPv1 or IGMPv2 membership message. The switch drops unknown multicast streams. The default is Enable.
   If disabled, the switch forwards IGMP Join/Leave PDUs, known multicast streams, and unknown multicast streams to the upstream mrouter interface.
8. Select the Fast Leave Auto-Assignment Mode **Enable** or **Disable** radio button.
   This selection enables or disables the automatic assignment of fast-leave messages to all ports and LAGs on the switch. The default is Enable.

   NOTE: For information about IGMP Plus mode, see Configure IGMP snooping automatically with IGMP Plus mode on page 255.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

11. To refresh the page, click the **Refresh** button.
    The following table describes the view-only fields on the page.

Table 81.IGMP snooping configuration information

Field	Description
Multicast Control Frame Count	The number of multicast control frames that are processed by the switch.
Interfaces Enabled for IGMP Snooping	The interfaces on which IGMP snooping is enabled.
Operational Mode	Indicates if IGMP snooping is globally enabled or disabled on the switch.
VLAN IDs Enabled For IGMP Snooping	The VLANs on which IGMP snooping is enabled.

Configure the IGMP snooping settings for interfaces

You can configure the IGMP snooping settings for interfaces.

To configure the IGMP snooping settings for interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > IGMP Snooping > Interface Configuration**.
   The Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.

4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Admin Mode** menu, select **Disable** or **Enable**.
   This selection enables or disables the administrative mode for IGMP snooping for the interface. The default is Disable.
6. In the **Membership Interval** field, specify the period that the switch waits for a group report before it removes the interface as a member of the group.
   Enter a value between 1 and 3600 seconds. The default is 600 seconds.
7. In the **Max Response Time** field, specify the period that the switch waits after it sent a query on an interface because it did not receive a report from a group on that interface.
   Enter a value that is 1 or greater but less than the value in the Membership Interval field. The default is 120 seconds.
8. In the **Expiration Time** field, specify the period that the switch waits to receive a query on the interface before it removes the interface from the list of interfaces with multicast routers attached.
   Enter a value from 0 to 3600 seconds. The default is 300 seconds. If you enter 0, you configure an infinite time-out (no expiration).
9. From the **Fast Leave** menu, select to enable or disable the IGMP snooping fast leave mode for the interface. This selection enables or disables the automatic assignment of fast-leave messages for the interface. The default is Disable.
   The Fast Leave Operational Mode field shows the status of the interface.
10. From the **Proxy Querier** menu, select to enable or disable the proxy querier for the interface. If disabled, the IGMP proxy query with source IP address 0.0.0.0 is not sent in response to an IGMP leave packet. The default is Enable.
11. Click the **Apply** button.
    Your settings are saved.
12. To save the settings to the running configuration, click the **Save** icon.

Configure IGMP snooping for VLANs automatically with IGMP Plus mode

IGMP Plus mode enables NETGEAR IGMP Plus™ to automatically configure IGMP snooping for VLANs, and build forwarding lists for multicast traffic.
You can also configure IGMP snooping for VLANs manually (see Configure IGMP snooping manually on page 256).

To configure IGMP snooping for VLANs automatically:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > IGMP Snooping > IGMP Snooping VLAN Configuration**.
   The IGMP Snooping VLAN Configuration page displays.
3. Select the check box next to the VLAN ID.
4. From the **IGMP Plus Mode** menu at the right of the page, select to enable or disable the IGMP Plus mode on the VLAN.
   If enabled, the following IGMP snooping modes are automatically enabled for the VLAN and the following occurs:
   - Admin mode
   - Fast-Leave
   - Proxy Querier
   - Report Flood Mode
   - Exclude Mrouter Interface Mode
   - Installs reserved Multicast MAC addresses into the system.

   If disabled, these IGMP snooping modes are automatically disabled for the VLAN.

   NOTE: For information about other settings on the page, see Configure IGMP snooping for VLANs manually on page 261.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Configure IGMP snooping for VLANs manually

You can manually configure the settings for IGMP snooping for VLANs, which is used to build forwarding lists for multicast traffic.
You can also configure IGMP snooping for VLANS automatically (see Configure IGMP snooping for VLANs automatically with IGMP Plus mode on page 260).

To configure the settings for IGMP snooping for a VLAN manually:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > IGMP Snooping > IGMP Snooping VLAN Configuration**.
   The IGMP Snooping VLAN Configuration page displays.
3. Select the check box next to the VLAN ID.
4. From the **Admin Mode** menu, select to enable or disable IGMP snooping for the VLAN.
   The default is Disable.
5. From the **Fast Leave** menu, select to enable or disable the IGMP snooping fast leave mode for the VLAN.
   This selection enables or disables the automatic assignment of fast-leave messages for all members of the VLAN. The default is Enable.
6. In the **Membership Interval** field, enter the period for the group membership interval of IGMP snooping for the VLAN.
   The period must be the value in the Maximum Response Time field plus a value from 1 to 3600 seconds. The default is 600 seconds.
7. In the **Maximum Response Time** field, enter the period for the maximum response time of IGMP snooping for the VLAN.
   The range must be from 1 to the value in the Group Membership Interval field minus 1. The value in the Maximum Response Time field must be greater than the value in the Group Membership Interval field. The default is 120 seconds.
8. In the **Multicast Router Expiry Time** field, enter the period for the multicast router expiration time of IGMP snooping for the VLAN.
   The range must be from 0 to 3600 seconds. The default is 300 seconds.

9. From the **Report Suppression** menu, select to enable or disable the IGMP snooping report suppression mode for the VLAN.
   This mode allows for the suppression of IGMP reports that are sent by multicast hosts. The switch does so by building a Layer 3 membership table and sending only the essential reports to IGMP routers that must receive the multicast traffic. As a result, the multicast report traffic that is sent to the IGMP routers is reduced. The default is Disable.
10. From the **Proxy Querier** menu, select to enable or disable the proxy querier for the VLAN.
    If disabled, the IGMP proxy query with source IP address 0.0.0.0 is not sent in response to an IGMP leave packet. The default is Enable.
11. From the **Report Flood Mode** menu, select to enable or disable the report flooding mode on the VLAN.
    If enabled, IGMP Join/Leave PDUs that the VLAN receives from a host on a downstream port are forwarded to all other ports in the VLAN. The default is Enable.
12. From the **Exclude Mrouter Interface Mode** menu, select to enable or disable the Exclude Mrouter Interface Mode.
    This selection specifies the type of information that is forwarded to the upstream multicast router interface:
    - **Enable**. The VLAN forwards IGMP Join/Leave PDUs that it receives on a downstream port to an upstream mrouter interface. In addition, the VLAN forwards a multicast data stream to an upstream mrouter interface only if that port already received an IGMPv1 or IGMPv2 membership message. The VLAN drops unknown multicast streams. The default is Enable.
    - **Disable**: The VLAN forwards IGMP Join/Leave PDUs, known multicast streams, and unknown multicast streams to the upstream mrouter interface.

    NOTE: For information about IGMP Plus mode for VLANs, see Configure IGMP snooping for VLANs automatically with IGMP Plus mode on page 260.
13. Click the **Apply** button.
    Your settings are saved.
14. To save the settings to the running configuration, click the **Save** icon.

Configure an IGMP multicast router interface

You can configure an interface as the designated interface to which a multicast router is attached. All IGMP packets snooped by the switch are forwarded to the multicast router that is reachable from this interface. We refer to this interface as the multicast router.
In most situations, this configuration is not required because the switch automatically detects a multicast router and forwards IGMP packets accordingly. This configuration might be required in a complex network if you want to make sure that the multicast router always receives IGMP packets from the switch.

To configure a multicast router interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > IGMP Snooping > Multicast Router Configuration**.
   The Multicast Router Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Multicast Router** menu, select to enable or disable the multicast router option.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Configure an IGMP multicast router VLAN

You can configure a specific VLAN for a specific interface to forward snooped IGMP packets to the multicast router that is connected to the interface.
In most situations, this configuration is not required because the switch automatically detects a multicast router and forwards IGMP packets accordingly. This configuration might be required in a complex network if you want to make sure that the multicast router always receives IGMP packets from the switch.

To configure a multicast router VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > IGMP Snooping > Multicast Router VLAN Configuration**.
   The Multicast Router VLAN Configuration page displays.
3. From the **Interface** menu, select the interface.
4. In the **VLAN ID** field, enter the VLAN ID.
5. From the **Multicast Router** menu, select **Enable** or **Disable** to specify if the VLAN is a multicast router VLAN.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

IGMP snooping querier overview

IGMP snooping requires that one central switch or router periodically query all end-devices on the network to announce their multicast memberships. This central device is the IGMP querier. The IGMP query responses, known as IGMP reports, keep the switch updated with the current multicast group membership on a port-by-port basis. If the switch does not receive updated membership information in a timely fashion, it stops forwarding multicasts to the port where the end device is located.
You can configure and display information for IGMP snooping queriers on the network and, separately, on VLANs.

Configure the IGMP snooping querier global settings

You can configure the global settings for an IGMP snooping querier on the switch.

To configure the global settings for an IGMP snooping querier:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > IGMP Snooping > Querier Configuration**.
   The Querier Configuration page displays.
3. Select the Querier Admin Mode **Enable** or **Disable** radio button to specify if IGMP snooping is enabled for the switch.
   The default is Disable.
4. In the **Snooping Querier IP Address** field, specify the snooping querier IP address that must be used as the source address in periodic IGMP queries.
   This address is used when no address is configured on the VLAN on which queries are sent.
5. In the **IGMP Version** field, specify the IGMP protocol version used in periodic IGMP queries.
   The version can be 1 to 2. The default is 2.
6. In the **Query Interval (secs)** field, specify the period in seconds between periodic queries sent by the snooping querier.
   The range is from 1 to 1800 seconds. The default is 60 seconds.
7. In the **Querier Expiry Interval (secs)** field, specify the period in seconds after which the last querier information is removed.
   The range is from 60 to 300 seconds. The default is 180 seconds.
8. Click the **Apply** button.
   Your settings are saved.
   The page displays the VLAN IDs enabled for IGMP snooping querier.

Configure an IGMP snooping querier for a VLAN

You can configure an IGMP querier for use with a VLAN on the network.

To configure an IGMP snooping querier for a VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > IGMP Snooping > Querier VLAN Configuration**.
   The Querier VLAN Configuration page displays.
3. Either specify a new VLAN or change an existing VLAN:
   - **Specify a new VLAN**: In the **VLAN ID** field, enter the ID for a new VLAN.
   - **Change an existing VLAN**: Select the check box for an existing VLAN.

4. From the **Querier Election Participate Mode** menu, select to enable or disable the querier election participate mode:
   - **Disabled**. If the switch detects another querier of the same version in the VLAN, the snooping querier moves to the non-querier state.
   - **Enabled**: The snooping querier participates in querier election, in which the lowest numbered IP address operates as the querier in that VLAN. The other querier moves to non-querier state.

   By default, the querier election participate mode is enabled for VLAN 1.
5. In the **Snooping Querier VLAN Address** field, specify the IP address to be used as the source IP address in periodic IGMP queries that are sent on the VLAN.
6. Either add the new VLAN or save the settings for the existing VLAN:
   - **Add the new VLAN**: Click the **Add** button. The VLAN is added.
   - **Save the settings for the existing VLAN**: Click the **Apply** button. Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 82.Querier VLAN configuration information

Field	Description
Operational State	The operational state of the IGMP snooping querier on a VLAN: • **Querier**: The snooping switch is the querier in the VLAN. The snooping switch sends out periodic queries with a time interval equal to the configured querier query interval. If the snooping switch finds a better querier in the VLAN, it moves to non-querier mode. • **Non-Querier**: The snooping switch is innNon-querier mode in the VLAN. If the querier expiry interval timer expires, the snooping switch moves into querier mode. • **Disabled**: The snooping querier is not operational on the VLAN. The snooping querier moves to disabled mode when IGMP snooping is not operational on the VLAN or when the querier address is not configured or the network management address is also not configured.
Operational Version	The operational IGMP protocol version of the querier.
Last Querier Address	The IP address of the last querier from which a query was snooped on the VLAN.
Last Querier Version	The IGMP protocol version of the last querier from which a query was snooped on the VLAN.
Operational Max Response Time	The maximum response time to be used in the queries that are sent by the snooping querier.

Remove the IGMP snooping querier settings for a VLAN

You can remove the IGMP snooping querier settings for a VLAN.

To remove the IGMP snooping querier settings for a VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Multicast > IGMP Snooping > Querier VLAN Configuration**.
   The Querier VLAN Configuration page displays.
3. Select the VLAN for which you want to remove the IGMP snooping querier settings.

4. Click the **Delete** button.
   Your settings are saved. (The VLAN itself is not deleted.)
5. To save the settings to the running configuration, click the **Save** icon.

Display the status of the IGMP snooping querier

You can display the status of the IGMP snooping querier for an interface or a VLAN.

To display the status of the IGMP snooping querier:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Multicast > IGMP Snooping Querier > IGMP Snooping Group Table**.
   The IGMP Snooping Group Table page displays.
3. To display information for a specific VLAN or interface, do the following:
   a. From the **Search** menu, select VLAN ID or interface.
   b. In the field, enter the VLAN ID or interface number.
   c. Click the **GO** button.
4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 83.IGMP snooping group information

Field	Description
VLAN ID	The ID of the VLAN on which the host sends IGMP member join requests.
Subscriber	The IP address and MAC address of the host that sends IGMP member join requests.
MC Group	The multicast IP address and MAC address that the subscriber is registered to.
Interface	The interface on which the IGMP member join requests us detected.
Type	The IGMP protocol version.
Timeout (secs)	The time that the most recent update for this host expires.

Multicast Listener Discovery snooping

In IPv6 networks, Multicast Listener Discovery (MLD) snooping performs a similar function as IGMP in IPv4 networks. With MLD snooping, IPv6 multicast data is selectively forwarded to ports that are configured to receive the data, instead of being flooded to all ports in a VLAN. The ports

are determined by snooping IPv6 multicast control packets.
A multicast listener is a device that is configured to receive IPv6 multicast packets. MLD is used by IPv6 multicast routers to discover the presence of multicast listeners on its directly-attached links and to discover which multicast packets are of interest to neighboring devices.
The MLD protocol is derived from IGMP. MLD version 1 (MLDv1) is equivalent to IGMPv2, and MLD version 2 (MLDv2) is equivalent to IGMPv3. MLD is a subprotocol of Internet Control Message Protocol version 6 (ICMPv6), and MLD messages are a subset of ICMPv6 messages, identified in IPv6 packets by a preceding Next Header value of 58.

Configure MLD snooping automatically with MLD Plus mode

MLD Plus mode lets you automatically configure MLD snooping, which is used to build forwarding lists for IPv6 multicast traffic.
You can also configure MLD snooping manually (see Configure MLD snooping manually on page 270).

To configure MLD snooping automatically:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > MLD Snooping > Configuration**.
   The Configuration page displays.
3. Select the MLD Plus Mode **Enable** or **Disable** radio button.
   If enabled, the following MLD snooping modes are automatically enabled:
   - MLD Snooping Admin mode
   - Exclude Mrouter Interface Mode

   The default is Enable.
   If disabled, these MLD snooping modes are automatically disabled.

   NOTE: For information about other settings on the page, see Configure MLD snooping manually on page 270.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

6. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 84.MLD snooping configuration information

Field	Definition
Multicast Control Frame Count	The number of multicast control frames that were processed.
Interfaces Enabled for MLD Snooping	The interface on which MLD snooping is administratively enabled. MLD snooping must be enabled globally and on an interface for the interface to be able to snoop MLD packets to determine which segments should receive multicast packets directed to the group address.
VLAN IDs Enabled For MLD Snooping	The VLAN on which MLD snooping is administratively enabled.

Configure MLD snooping manually

You can manually configure the settings for MLD snooping, which is used to build forwarding lists for IPv6 multicast traffic.
You can also configure MLD snooping automatically (see Configure MLD snooping automatically with MLD Plus mode on page 269).

To configure the settings for MLD snooping manually:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > MLD Snooping > Configuration**.
   The Configuration page displays.
3. Select the MLD Snooping Admin Mode **Enable** or **Disable** radio button to specify the administrative mode for MLD snooping for the switch. The default is Disable.
4. Select the Proxy Querier Mode **Enable** or **Disable** radio button.
   This enables or disables an MLD proxy querier on the system. If it is disabled, then an MLD proxy query with source IP 0::0 is not sent in response to an MLD leave packet. If it is enabled, then MLD proxy queries are sent. The default is Enable.

5. Select the Exclude Mrouter Interface Mode **Enable** or **Disable** radio button.
   This selection specifies the type of information that is forwarded to the upstream multicast router interface.
   If enabled, the switch blocks all unknown multicast data through the mrouter port, whether the port is configured dynamically or statically. Only MLD PDUs are allowed to pass through the mrouter port to the upstream router interface.
   The default is Enable. If disabled, the switch forwards both unknown multicast data and MLD PDUs to the upstream multicast router interface.

   NOTE: For information about MLD Plus mode, see Configure MLD snooping automatically with MLD Plus mode on page 269.

6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.
8. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 85.MLD snooping configuration information

Field	Definition
Multicast Control Frame Count	The number of multicast control frames that were processed.
Interfaces Enabled for MLD Snooping	The interface on which MLD snooping is administratively enabled. MLD snooping must be enabled globally and on an interface for the interface to be able to snoop MLD packets to determine which segments should receive multicast packets directed to the group address.
VLAN IDs Enabled For MLD Snooping	The VLAN on which MLD snooping is administratively enabled.

Configure the MLD snooping settings for interfaces

You can configure the MLD snooping settings for interfaces.

To configure the MLD snooping settings for interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Switching > Multicast > MLD Snooping > Interface Configuration**.
   The Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Admin Mode** menu, select to enable or disable MLD snooping for the interface. The default is Disable.
6. In the **Membership Interval** field, specify the time that the switch must wait for a report for a particular group on a particular interface before it deletes that interface from the group. The valid range is from 2 to 3600 seconds. The configured value must be greater than the maximum response time. The default is 260 seconds.
7. In the **Max Response Time** field, specify the time that the switch must wait after sending a query on an interface because it did not receive a report for a particular group on that interface.
   Enter a value greater than or equal to 1 and less than the membership interval in seconds. The default is 10 seconds. The configured value must be less than the membership interval.
8. In the **Expiration Time** field, specify the time that the switch must wait to receive a query on an interface before removing the interface from the list of interfaces with multicast routers attached.
   Enter a value between 0 and 3600 seconds. The default is 300 seconds. A value of zero indicates an infinite time-out, that is, no expiration.

9. From the **Fast Leave** menu, select **Enable** or **Disable** to enable or disable Fast Leave on the interface.
   If Fast Leave is enabled, the interface can immediately be removed from the Layer 2 forwarding table when the switch receives an MLD leave message for a multicast group without first sending MAC-based general queries. The default is Disable.
10. From the **Proxy Querier** menu, select **Enable** or **Disable** to enable or disable the proxy querier mode on the interface.
    If the mode is disabled, an MLD proxy query with source IP 0::0 is not sent in response to an MLD leave packet. If the mode is enabled, MLD proxy queries are sent. The default value is Enable.
11. Click the **Apply** button.
    Your settings are saved.
12. To save the settings to the running configuration, click the **Save** icon.

Configure MLD snooping for VLANs automatically with MLD Plus mode

MLD Plus mode lets you automatically configure MLD snooping for VLANs, which is used to build forwarding lists for IPV6 multicast traffic.
You can also configure MLD snooping for VLANs manually (see Configure IGMP snooping for VLANs manually on page 261).

To configure MLD snooping for VLANs automatically:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > MLD Snooping > MLD VLAN Configuration**.
   The MLD VLAN Configuration page displays.
3. Select the check box for the VLAN ID for which MLD snooping must be enabled.

4. From the **MLD Plus Mode** menu, select to enable or disable the MLD Plus mode on the VLAN.
   If enabled, the following MLD snooping modes are automatically enabled for the VLAN:
   - Admin mode
   - Fast-Leave
   - Exclude Mrouter Interface Mode

   If disabled, these MLD snooping modes are automatically disabled for the VLAN.

   NOTE: For information about other settings on the page, see Configure IGMP snooping for VLANs manually on page 261.

5. Click **Add** to enable MLD Snooping on the specified VLAN.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Configure MLD snooping for VLANs manually

You can manually configure the settings for MLD snooping for VLANs, which is used to build forwarding lists for IPv6 multicast traffic.
You can also configure MLD snooping for VLANS automatically (see Configure MLD snooping for VLANs automatically with MLD Plus mode on page 273).

To configure the settings for MLD snooping for a VLAN manually:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > MLD Snooping > MLD VLAN Configuration**.
   The MLD VLAN Configuration page displays.
3. Either specify a new VLAN or change an existing VLAN:
   - **Specify a new VLAN**: In the VLAN ID field, enter the ID for a new VLAN.
   - **Change an existing VLAN**: Select the check box for an existing VLAN.
4. From the **Fast Leave** menu, select to enable or disable the MLD snooping Fast Leave mode.
5. In the **Membership Interval** field, specify the period for the group membership interval of MLD snooping.
   The range is from the period that you specify in the **Maximum Response Time** field *plus 1* to 3600 seconds. The default is 260 seconds.

6. In the **Maximum Response Time** field, specify the period for the maximum response time of MLD snooping.
   The range is from 1 second to the period that you specify in the **Membership Interval***minus 1*.
   The default is 10 seconds.
7. In the **Multicast Router Expiry Time** field, specify the period for the multicast router expiration time of MLD Snooping.
   The range is from 0 to 3600 seconds. The default is 300 seconds.
8. From the **Proxy Querier Mode** menu, select to enable or disable the proxy querier mode.
   If you select Disable, the MLD proxy query with source IP 0::0 is not sent in response to an MLD leave packet. The default value is Enable.
9. From the **Exclude Mrouter Interface Mode** menu, select to enable or disable the mrouter interface mode.
   This selection specifies the type of information that is forwarded to the upstream multicast router interface.
   If enabled, the interface blocks all unknown multicast data through the mrouter port, whether the port is configured dynamically or statically. Only MLD PDUs are allowed to pass through the mrouter port to the upstream router interface.
   The default is Enable. If disabled, the interface forwards both unknown multicast data and MLD PDUs to the upstream multicast router interface.

   NOTE: For information about MLD Plus mode for VLANs, see Configure MLD snooping for VLANs automatically with MLD Plus mode on page 273.
10. Either add the new VLAN or save the settings for the existing VLAN:
    - **Add the new VLAN**: Click the **Add** button. The VLAN is added.
    - **Save the settings for the existing VLAN**: Click the **Apply** button. Your settings are saved.
11. To save the settings to the running configuration, click the **Save** icon.

Remove the MLD snooping querier settings for a VLAN

You can remove the MLD snooping querier settings for a VLAN.

To remove the MLD snooping querier settings for a VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Multicast > MLD Snooping > Querier VLAN Configuration**.
   The Querier VLAN Configuration page displays.

3. Select the VLAN for which you want to remove the IGMP snooping querier settings.
4. Click the **Delete** button.
   Your settings are saved. (The VLAN itself is not deleted.)
5. To save the settings to the running configuration, click the **Save** icon.

Configure an MLD multicast router interface

You can configure an interface as the designated interface to which a multicast router is attached. All MLD packets snooped by the switch are forwarded to the multicast router that is reachable from this interface. We refer to this interface as the multicast router.
In most situations, this configuration is not required because the switch automatically detects a multicast router and forwards MLD packets accordingly. This configuration might be required in a complex network if you want to make sure that the multicast router always receives MLD packets from the switch.

To configure an MLD multicast router interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > MLD Snooping > Multicast Router Configuration**.
   The Multicast Router Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

5. From the **Multicast Router** menu, select to enable or disable the multicast router option.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Configure an MLD multicast router VLAN

You can configure a specific VLAN for a specific interface to forward snooped MLD packets to the multicast router that is connected to the interface.
In most situations, this configuration is not required because the switch automatically detects a multicast router and forwards MLD packets accordingly. This configuration might be required in a complex network if you want to make sure that the multicast router always receives MLD packets from the switch.

To configure a multicast router VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > MLD Snooping > Multicast Router VLAN Configuration**.
   The Multicast Router VLAN Configuration page displays.
3. From the **Interface** menu, select the interface.
4. In the **VLAN ID** field, enter the VLAN ID.
5. From the **Multicast Router** menu, select **Enable** or **Disable** to specify if the VLAN is a multicast router VLAN.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Configure the MLD snooping querier global settings

You can configure the global settings for an MLD snooping querier on the switch.

To configure the global settings for an MLD snooping querier:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Switching> Multicast > MLD Snooping > Querier Configuration**.
   The Querier Configuration page displays.
3. Select the Querier Admin Mode **Enable** or **Disable** radio button to specify if MLD snooping is enabled for the switch.
   The default is Disable.
4. In the **Querier Address** field, specify the snooping querier IP address that must be used as the source address in periodic MLD queries.
   This address is used when no address is configured on the VLAN on which queries are sent.
   The supported IPv6 formats are x:x:x:x:x:x:x:x and x::x.

   NOTE: The MLD Version field always states MLD protocol version 1. This is the version that is used in periodic MLD queries.
5. In the **Query Interval (secs)** field, specify the period in seconds between periodic queries sent by the snooping querier.
   The range is from 1 to 1800 seconds. The default is 60 seconds.
6. In the **Querier Expiry Interval (secs)** field, specify the period in seconds after which the last querier information is removed.
   The range is from 60 to 300 seconds. The default is 60 seconds.
7. Click the **Apply** button.
   Your settings are saved.
   The page displays the VLAN IDs enabled for MLD snooping querier.
8. To save the settings to the running configuration, click the **Save** icon.

Configure an MLD snooping querier for a VLAN

You can configure an MLD querier for use with a VLAN on the network.

To configure an MLD snooping querier for a VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > MLD Snooping > Querier VLAN Configuration**.
   The Querier VLAN Configuration page displays.
3. Either specify a new VLAN or change an existing VLAN:
   - **Specify a new VLAN**: In the VLAN ID field, enter the ID for a new VLAN.
   - **Change an existing VLAN**: Select the check box for an existing VLAN.

4. From the **Querier Election Participate Mode** menu, select to enable or disable the querier election participate mode:
   - **Disable**: If the switch detects another querier of the same version in the VLAN, the snooping querier moves to the non-querier state.
   - **Enable**: The snooping querier participates in querier election, in which the lowest numbered IPv6 address operates as the querier in that VLAN. The other querier moves to non-querier state.
5. In the **Querier VLAN Address** field, specify the IPv6 address to be used as the source address in periodic IGMP queries that are sent on the VLAN.
6. Either add the new VLAN or save the settings for the existing VLAN:
   - **Add the new VLAN**: Click the **Add** button. The VLAN is added.
   - **Save the settings for the existing VLAN**: Click the **Apply** button. Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 86.Querier VLAN configuration information

Field	Description
Operational State	The operational state of the MLD snooping querier on a VLAN: • **Querier**: The snooping switch is the querier in the VLAN. The snooping switch sends out periodic queries with a time interval equal to the configured querier query interval. If the snooping switch finds a better querier in the VLAN, it moves to non-querier mode. • **Non-Querier**: The snooping switch is in non-querier mode in the VLAN. If the querier expiry interval timer expires, the snooping switch moves into querier mode. • **Disabled**: The snooping querier is not operational on the VLAN. The snooping querier moves to disabled mode when MLD snooping is not operational on the VLAN or when the querier address is not configured or the network management address is also not configured.
Operational Version	The operational MLD protocol version of the querier
Last Querier Address	The IP address of the last querier from which a query was snooped on the VLAN
Last Querier Version	The MLD protocol version of the last querier from which a query was snooped on the VLAN
Operational Max Response Time	The maximum response time to be used in the queries that are sent by the snooping querier

Remove the MLD snooping querier settings for a VLAN

You can remove the MLD snooping querier settings for a VLAN.

To remove the MLD snooping querier settings for a VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> Multicast > MLD Snooping > Querier VLAN Configuration**.
   The Querier VLAN Configuration page displays.
3. Select the VLAN for which you want to remove the IGMP snooping querier settings.

4. Click the **Delete** button.
   Your settings are saved. (The VLAN itself is not deleted.)
5. To save the settings to the running configuration, click the **Save** icon.

Multicast VLAN registration

IGMP and MLD snooping help to limit multicast traffic when member ports are in the same VLAN. However, when ports belong to different VLANs, a copy of the multicast stream is sent to each VLAN with member ports in the multicast group. Multicast VLAN registration (MVR) eliminates the need to duplicate the multicast traffic when multicast group member ports belong to different VLANs.
MVR uses a dedicated multicast VLAN to forward multicast traffic over the L2 network. You can configure one multicast source VLAN (MVLAN) only on the switch. Such an MVLAN is used only for certain multicast traffic, such as traffic from an IPTV application, to prevent duplication of multicast streams for clients in different VLANs. Clients can dynamically join or leave the MVLAN without interfering with their membership in other VLANs.
MVR, like IGMP and MLD snooping, allows a Layer 2 switch to listen to IGMP and MLD messages to learn about multicast group membership.
You can configure global, group, interface, and group membership settings.

Configure the global MVR settings

You can configure the global MVR settings that apply to the switch.

To configure the global MVR settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching>MVR>Basic >MVR Configuration**.
   The MVR Configuration page displays.
3. From the **MVR Running** menu, select **Enable** or **Disable** to specify if MVR is globally enabled or disabled on the switch
   The default is Disable.
4. In the **MVR Multicast VLAN** field, specify the VLAN on which MVR multicast data is received. All source ports belong to this VLAN. The range is from 1 to 4093. The default 1.

5. In the **MVR Global Query Response Time** field, set the period that the switch must wait for an IGMP group membership report from an interface before removing the interface as a member from the multicast group.
   This period applies only to the removal of the interface from the receiver port on the switch. When an IGMP query is sent from a receiver port, the switch waits for the default or configured MVR query time for an IGMP group membership report before removing the interface as a member from the multicast group. The period is equal to a tenths of a second. The range is from 1 to 100 tenths. The default is 5 tenths or one-half.
6. Select the MVR Mode **compatible** or **dynamic** radio button to specify the MVR mode of operation.
   Select the **compatible** radio button (the default mode) to block IGMP group membership reports on source ports. Select the **dynamic** radio button to allow IGMP group membership reports on source ports.
7. Click the **Apply** button.
   Your settings are saved.
8. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 87.MVR configuration information

Field	Definition
MVR Max Multicast Groups	The maximum number of multicast groups that MVR supports.
MVR Current Multicast Groups	The number of the MVR groups allocated.

Configure an MVR group

You can configure an MVR group. After you configure interfaces for MVR, you can add them as members to the MVR group.

To configure an MVR group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching>MVR>Advanced >MVR Group Configuration**.
   The MVR Group Configuration page displays.
3. In the **MVR Group IP** field, specify the IP address for the new MVR group.

4. In the **Count** field, specify the number of contiguous MVR groups.
   This number helps you to create multiple MVR groups through a single click of the **Add** button. If the field is empty, then clicking the button creates only one new group. The field is displayed as empty for each particular group. The range is from 1 to 256.
5. Click the **Add** button.
   The MVR group is added.
6. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 88.MVR group configuration information

Field	Definition
Status	The status of the specific MVR group.
Members	The list of interfaces that participate in the MVR group.

Remove an MVR group

You can remove an MVR group that you no longer need.

To remove an MVR group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > MVR > Advanced > MVR Group Configuration**.
   The MVR Group Configuration page displays.
3. Select the check box for the MVR group.
4. Click the **Delete** button.
   The MVR group is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure an MVR interface

We recommend that you first configure an MVR interface before you add it as a member to an MVR group.

To configure an MVR interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching>MVR>Advanced >MVR Interface Configuration**.
   The MVR Interface Configuration page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Admin Mode** menu, specify whether MVR is enabled on the interface by selecting **Enable** or **Disable**.
   The default is Disable.
6. From the **Type** menu, specify whether the interface is an MVR receiver or an MVR source by selecting **receiver** or **source**.
   The default port type is none.
7. From the **Immediate Leave** menu, specify whether the Immediate Leave feature is enabled by selecting **Enable** or **Disable**.
   The default is Disable.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.
10. To refresh the page, click the **Refresh** button.
    The Status field shows the status for each interface.

Configure the interface members of an MVR group

You can add or remove interfaces as members of an MVR group.

To configure the interface members of an MVR group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching>MVR>Advanced >MVR Group Membership**.
   The MVR Group Membership page display.
   The ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed, and you can select ports on different stacked switches.
3. From the **Group IP** menu, select the IP address of the MVR group.
4. In the Ports table (or if a stack is configured, in one or more of the Ports tables), click each port that you want to make a member of the MVR group.
   A selected port is shown by a check mark.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Display the MVR statistics

You can display MVR statistics for the switch. These statistics are associated with IGMP.

To display the MVR statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching>MVR>Advanced >MVR Statistics**.
   The MVR Statistics page displays.

3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 89.MVR statistics information

Field	Definition
IGMP Query Received	The number of received IGMP queries.
IGMP Report V1 Received	The number of received IGMP V1 reports.
IGMP Report V2 Received	The number of received IGMP V2 reports.
IGMP Leave Received	The number of received IGMP leaves.
IGMP Query Transmitted	The number of transmitted IGMP queries.
IGMP Report V1 Transmitted	The number of transmitted IGMP V1 reports.
IGMP Report V2 Transmitted	The number of transmitted IGMP V2 reports.
IGMP Leave Transmitted	The number of transmitted IGMP leaves.
IGMP Packet Receive Failures	The number of IGMP packet receive failures.
IGMP Packet Transmit Failures	The number of IGMP packet transmit failures.

MAC address table

You can view or configure the MAC address table. This table contains information about unicast entries for which the switch holds forwarding or filtering information. This information lets the transparent bridging function determine how an incoming frame must be propagated.

View, search, or clear the MAC address table

If you clear the MAC address entries in the MAC address table, only the dynamic entries are removed.

To view, search, or clear the MAC address table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Address Table > Basic > Address Table**.
   The Address Table page displays.

3. Use the **Search** menu and field to search for a MAC address, VLAN ID, or interface number:
   - **Search by MAC Address**: From the **Search** menu, select MAC Address, and enter the 6-byte hexadecimal MAC address in two-digit groups separated by colons, for example, 01:23:45:67:89:AB. Then click the **Go** button. If the address exists, that entry is displayed as the first entry followed by the remaining (higher) MAC addresses. An exact match is required.
   - **Search VLAN ID**: From the **Search** menu, select VLAN ID, and enter the VLAN ID, for example, 100. Then click the **Go** button.
   - **Search Interface**: From the **Search** menu, select **Port**, and enter the interface ID using the respective interface naming convention (for example, 0/2). Then click the **Go** button.
4. To remove the dynamic entries from the MAC address table, click the **Clear** button.
5. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 90.MAC address table information

Field	Description
Total MAC Address	The number of total MAC addresses learned or configured.
MAC Address	The unicast MAC address for which the switch holds forwarding information, filtering information, or both forwarding and filtering information. The format is a 6-byte MAC address that is separated by colons, for example 01:23:45:67:89:AB.
VLAN ID	The VLAN ID associated with the MAC address.
Port	The interface on which the address was learned.
Status	The status of this entry: ○ **Static**: The MAC address was added by the switch or a user and cannot be relearned. ○ **Learned**: The MAC address was learned, and is being used. ○ **Management**: The management MAC address.

Set the dynamic address aging interval

You can set the address aging interval for the forwarding database. This is the time-out period in seconds for aging out dynamically learned forwarding information.

To set the dynamic address aging interval:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Address Table > Advanced > Dynamic Addresses**.
   The Dynamic Addresses page displays.
3. In the **Address Aging Timeout (seconds)** field, specify the time-out period in seconds for aging out dynamically learned forwarding information.
   The period is from 10 to 1000000 seconds. The default is 3000 seconds.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Add a static MAC address to the MAC address table

Static MAC address entries are the ones that you manually add to the MAC address table for a specific interface and VLAN.

To add a static MAC address to the MAC address table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Address Table > Advanced > Static MAC Address**.
   The Static MAC Address page displays.
3. From the **Interface** menu, select the interface or LAG.
4. In the **Static MAC Address** field, enter the MAC address.
5. From the **VLAN ID** menu, select the VLAN ID that must be associated with the MAC address.
6. Click the **Add** button.
   The static MAC address is added to the MAC address table.
7. To save the settings to the running configuration, click the **Save** icon.

Remove a static MAC address from the MAC address table

You can remove a static MAC address that you no longer need.

To remove a static MAC address from the MAC address table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Address Table > Advanced > Static MAC Address**.
   The Static MAC Address page displays.
3. From the **Interface** menu, select the interface or LAG.
4. Select the check box for the MAC address.
5. Click the **Delete** button.
   The static MAC address is removed from the MAC address table.
6. To save the settings to the running configuration, click the **Save** icon.

Port settings

For the physical ports and LAGs on the switch, you can configure and display the settings, including the administrative mode of a port or LAG (by default, enabled) and speed settings of a port (by default, automatically detected), add port descriptions, display information about optional port transceivers that you might have installed, and configure the port link flap settings.

Configure and display the port settings

You can configure and display the information for the physical port on the switch. Some settings that apply to physical ports do not apply to LAGs.

NOTE: If you change the autonegotiation, speed, or duplex mode for a physical port, the switch might be inaccessible for a number of seconds while the new settings take effect.

To configure and display the port settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Ports > Port Configuration**.
   The Port Configuration page display.

3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **STP Mode** menu, select to enable or disable the Spanning Tree Protocol administrative mode for the port or LAG.
   The default is Enable.
6. From the **Admin Mode** menu, select to enable or disable the administrative mode for the port or LAG.
   For the port or LAG to participate in the network, you must select **Enable**. The default is Enable.
7. From the **LACP Mode** menu, select to enable or disable the Link Aggregation Control Protocol administrative mode for the port.
   The mode must be enabled for the port to participate in link aggregation. The default is Enable.
8. From the **Autonegotiation** menu, select to enable or disable the speed autonegotiation mode for the port.
   The default is Enable.

9. To change the speed for the port, in the **Speed** field, type the supported speed:
   For Gigabit Ethernet ports:
   - **Auto**: The speed is set by the auto-negotiation process. This is the default setting.
   - **1000**: The speed is set to 1000 Mbits/second.
   - **100**: The speed is limited to 100 Mbits/second.

   For multispeed Ethernet ports:
   - **Auto**: The speed is set by the auto-negotiation process. This is the default setting.
   - **10G**: The speed is set to 10 Gbits/second.
   - **5G**: The speed is limited to 5 Gbits/second.
   - **2.5G**: The speed is limited to 2.5 Gbits/second.
   - **1000**: The speed is limited to 1000 Mbits/second.

   The delimiter characters for setting different speed values are a comma (,), a period (.) and a space ( ). For you to set the auto-negotiation speed, the auto-negotiation mode selection must be **Enable**.

   NOTE: You can set the speed for Ethernet ports. For SFP ports, the speed is automatically detected and you cannot change it.

10. To change the duplex mode for the port, from the **Duplex Mode** menu, select one of the following values are as follows:
    - **Auto**: The duplex mode is set by the auto-negotiation process. This is the default setting.
    - **Full**: Transmission between the devices occurs in both directions simultaneously.
    - **Half**: Transmission between the devices occurs in only one direction at a time.

11. From the **Link Trap** menu, select to enable or disable the option to send a trap when the port or LAG link status changes.
    For ports, the default is Enable. For LAGs, the default is Disable.

12. In the **Frame Size** field, specify the maximum Ethernet frame size that the port or LAG supports, including the Ethernet header, CRC, and payload.
    The range is from 1500 to 12270. The default is 9198.

13. In the **Debounce Time** field, specify the wait period for port or LAG debouncing.
    Specify the period as a multiple of 100 milliseconds (msec) in the range from 100 to 5000.
    The default is 0 seconds, which means that debouncing is disabled.

14. From the **Flow Control** menu, select the configuration for IEEE 802.3x flow control for the port:
    - **Disable**: If the port buffers become full, the switch does not send pause frames, and data loss could occur. This is the default setting.
    - **Symmetric**: If the port buffers become full, the switch sends pause frames to stop traffic. Flow control helps to prevent data loss when the port cannot keep up with the number of frames being switched. When you enable flow control, the switch can send a pause frame to stop traffic on the port if the amount of memory used by the packets on the port exceeds a preconfigured threshold and responds to pause requests from partner devices. The paused port does not forward packets for the time that is specified in the pause frame. When the pause frame time elapses, or the utilization returns to a specified low threshold, the switch enables the port to again transmit frames. The switch also honors incoming pause frames by temporarily halting transmission.
    - **Asymmetric**: If the port buffers become full, the switch does not send pause frames, and data loss could occur. However, the switch does honor incoming pause frames by temporarily halting transmission.
15. In the **Load Interval** field, specify the load interval period for the port or LAG.
    This is the period for which data is used to compute load statistics. Enter the period in multiples of 30 seconds from 30 to 600. The default is 300 seconds. The smaller the load interval is, the more accurate the instantaneous rate for load statistics is.
16. Click the **Apply** button.
    Your settings are saved.
17. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 91.Port configuration information

Field	Description
Media Type	The media type that the port supports.
Port Type	For normal ports this field is Normal. Otherwise the possible values are as follows: • **Mirrored**: The port is a mirrored port on which all the traffic is copied to the probe port. • **Probe**: Use this port to monitor a mirrored port. • **Trunk Member**: The port is a member of a link aggregation trunk. Look at the LAG pages for more information.
Physical Status	The port speed and duplex mode.
Link Status	Indicates whether the port or LAG link is up or down.
ifIndex	The ifIndex of the interface table entry associated with the port or LAG.

Add port, LAG, and VLAN descriptions

You can add a description for a port, LAG, and VLAN.

To add port, LAG, and VLAN descriptions:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Ports > Port Description**.
   The Port Description page displays.

3. Select whether to display physical interfaces, LAGs, VLANs, or all by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Physical interfaces, LAGs, and VLANs are displayed, or for a switch stack, physical interfaces on all switches in the stack, LAGs, and VLANs are displayed.
4. Select one or more interfaces or VLANs by taking one of the following actions:
   - To configure a single interface or VLAN, select the check box associated with the interface or VLAN, or type the interface or VLAN number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces or VLANs with the same settings, select the check box associated with each interface or VLAN.
   - To configure all interfaces, VLANs, or both interfaces and VLANs with the same settings, select the check box in the heading row.
5. In the **Description** field, enter a description of up to 64 characters.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 92.Port description information

Field	Description
MAC Address	The MAC address of the interface or VLAN.
PortList Bit Offset	The bit offset value that corresponds to the interface or VLAN when the MIB object type PortList is used to manage in SNMP.
ifIndex	The interface index associated with the interface or VLAN.

Display transceiver module information

To enable high-speed fiber and Gigabit Ethernet, short- and long-distance connections on the switch, SFP, SFP+, and SFP28 fiber ports can accommodate standard 1G SFP, 10G SFP+, and SFP28 25G transceiver modules. You can view the transceiver module information for all fiber ports on the switch.

To view transceiver module information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Ports > Port Transceiver**.
   The Port Transceiver page displays.
3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 93.Port transceiver information

Field	Description
Port	The port in which the transceiver module is installed.
Vendor Name	The vendor name of the transceiver module.
Link Length 50 µm [m]	The link length supported for 50 µm fiber.
Link Length 62.5 µm [m]	The link length supported for 62.5 µm fiber.
Serial Number	The serial number of the transceiver module.
Part Number	The part number of the transceiver module.
Nominal Bit Rate [Mpbs]	The nominal signalling rate for the transceiver module.
Revision	The vendor revision of the transceiver module.
Compliance	The compliance of the transceiver module.
Supported	Indicates if the inserted transceiver module is actually supported (Yes or No).
Possible Speed Detected	The potential speed that the transceiver module can support.

Configure the expandable port settings

You can manage the expandable port configurations for M4350 switches with 100GBASE-X QSFP28 ports. Expandable port configuration allows the 100G QSFP28 ports on M4350 switches to operate in breakout mode, and to function as:

- One 100G interface.
- Two 50G interfaces.
- Four 25G interfaces.

Configuring a 100G port in breakout mode enables connectivity to multiple devices using a single QSFP28 uplink port. If you configure all supported 100G ports in breakout mode, the switch provide up to 16 additional 25G interfaces (model dependent). You can manage the speed settings and configure the expanded interfaces on the Port Configuration page. For more information, see Configure and display the port settings on page 289.

To configure the expandable port settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Ports > Expandable Port Configuration**.
   The Expandable Port Configuration page displays.
   a. Select the 100G port you want to configure.
   b. Change the Configured Mode to one of the following: 4x25G.
   c. Click Apply.
3. Select the 100 Gbps port to be configured.
4. Change the Configured Mode to one of the following:
   - 4X25G
   - 2X50G
   - 1X100G
5. Click the **Apply** button.
   Your settings are saved.
   If you selected 2X50G or 4X25G, the Expandable Port Configuration page displays the operational mode of the port and the port numbers of the expanded interfaces. You can view and manage the expanded ports on the Port Configuration page.

Configure the port link flap settings

You can configure the port link flap settings, which determine when a port is automatically placed in the disabled state. You can also configure the automatic recovery settings, which allows a port to be automatically activated again.

To configure the port link flap settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > Ports > Link Flap Configuration**.
   The Link Flap Configuration page displays.
3. Select the Admin Mode **Enable** or **Disable** radio button to enable or disable the link flap administrative mode.
   For you to be able to configure the link flap settings, you must select Enable. The default is Disable.
4. In the **Max-Count** field, enter the maximum number of flaps that are allowed before the port is placed in the disabled state.
   You can enter a number from 2 to 10. By default, the number is 5.
5. In the **Duration** field, enter the maximum period in seconds during which the number of link flaps is counted.
   If the number of link flaps on the port is greater than or equal to the number that you enter in the **Max-Count** field, the port is placed in the disabled state. By default, the period is 10 seconds. You can enter a period from 3 to 200 seconds.
6. Select the Auto-Recovery Admin Mode **Enable** or **Disable** radio button to enable or disable the auto-recovery administrative mode.
   To enable auto-recovery, you must select **Enable**. If enabled, the port is automatically removed from the disabled state after the interval is reached. The default is Disable.
7. In the **Auto-Recovery Interval** field, enter the period in seconds after which the port is automatically removed from the disabled state.
   You can enter a period from 30 to 86400 seconds. The default period is 300 seconds.
8. Click the **Apply** button.
   Your settings are saved.
   The D-Disabled Ports due to Link Flap section lists the ports that are in the disabled state.
9. To save the settings to the running configuration, click the **Save** icon.

Link aggregation groups

Link aggregation groups (LAGs), which are also known as port channels, allow you to combine multiple full-duplex Ethernet links into a single logical link. Network devices treat the aggregation as if it were a single link, which increases fault tolerance and provides load sharing. You assign the LAG VLAN membership after you create a LAG. The LAG by default becomes a member of the management VLAN.

A LAG interface can be either static or dynamic, but not both. All members of a LAG must participate in the same protocols. A static port-channel interface does not require a partner system to be able to aggregate its member ports.

The switch supports static LAGs. When a port is added as a static member to a LAG, the port neither transmits nor receives LACPDUs.

Auto-LAG overview

An Auto-LAG is a LAG that forms automatically between two devices that support the Auto-LAG feature. An Auto-LAG is a dynamic Layer 2 LAG that is based on the Link Aggregation Control Protocol (LACP).

NOTE: A LAG is also referred to as a port channel or an EtherChannel.

The switch can detect the physical links with a partner device and automatically configure a LAG (that is, an Auto-LAG) on interconnected and capable ports at both ends. The switch can form one Auto-LAG only with each partner device.

The Auto-LAG feature functions together with the Auto-Trunk feature, which must also be supported and enabled on the partner device with which the LAG is formed. After an Auto-LAG is formed, the switch automatically applies trunk mode (that is, an Auto-Trunk) to the LAG at both ends. In other words, after an Auto-LAG is formed, the mode for the ports that participate in an Auto-LAG changes from the default switch port mode to the trunk port mode. For more information about the Auto-Trunk feature, see Auto-Trunk overview on page 199.

For the switch to form an Auto-LAG with a partner device, the following are required:

- Both the Auto-LAG and Auto-Trunk features must be supported and globally enabled on the switch and the partner device. (On the NETGEAR switch, the Auto-LAG and Auto-Trunk features are enabled by default.)
- At least two links must be established between the switch and the partner device, and these links must support the same speed and duplex mode.
- The links cannot be members of a manually configured static or dynamic LAG.
- LLDP must be enabled on the interconnected ports on the switch and the partner device. (On the NETGEAR switch, LLDP is enabled by default on all ports.)
- The interconnected ports on the switch and the partner device must be in the default switch

port mode, which is the General mode. If the ports are in the Access mode or already in the Trunk mode, an Auto-Trunk cannot be formed on the Auto-LAG.

An Auto-LAG can form with up to eight interfaces as members. Interfaces are automatically selected for the Auto-LAG based on whether they are up and available and on the following conditions:

- The interface is not already manually configured as a member of a LAG.
- The interface is not manually configured as a trunk port or an access port. That is, the interface must be a general interface.

NOTE: The switch can support multiple static and dynamic LAGs, but with each partner device, the switch can support a single Auto-LAG only.

Enable or disable Auto-LAGs

By default, the Auto-LAG feature is globally enabled but you can globally disable it.

To enable or disable Auto-LAGs:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> LAG > LAG Configuration**.
   The LAG Configuration page displays.
3. In the LAG Global Configuration section, select the Auto-LAG Admin Mode **Enable** or **Disable** radio button.
   By default, the Auto-LAG feature is globally enabled.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Enable or disable the Auto-LAG global AVB profile mode

The Audio Video AVB profile is one of the network profiles that is supported in the NETGEAR Engage controller. This network profile can be used for IP audio devices that support Audio Video Bridging (AVB). By default, the Auto-LAG global AVB profile mode is globally enabled but you can globally disable it.

To enable or disable the Auto-LAG global AVB profile mode:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> LAG > LAG Configuration**.
   The LAG Configuration page displays.
3. In the LAG Global Configuration section, select the Auto-LAG Global AVB Profile Mode **Enable** or **Disable** radio button.
   By default, the Auto-LAG global AVB profile mode is globally enabled.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the hash mode for Auto-LAGs

By default, the Auto-LAG feature is enabled and uses a hash mode that auto-configures a LAG based on the destination MAC address, VLAN, EtherType, and incoming port in the packet (that is, the hash mode with the name 2 Dest MAC, VLAN, EType, incoming port). You can change the hash mode (that is, the load balancing mode) for the Auto-LAG feature.
The switch balances traffic on a LAG by selecting one of the links in the channel over which packets must be transmitted. The switch selects the link by creating a binary pattern from selected fields in a packet and associating that pattern with a particular link. The hash mode determines which fields in a packet the switch selects.

To change the hash mode for the Auto-LAGs:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> LAG > LAG Configuration**.
   The LAG Configuration page displays.

3. In the LAG Global Configuration section, from the **Auto-LAG Global Hash Mode** menu, select the hash mode for the Auto-LAGs:
   - **1 Src MAC, VLAN, EType, incoming port**: Source MAC address, VLAN, EtherType, and incoming port associated with the packet.
   - **2 Dest MAC, VLAN, EType, incoming port**: Destination MAC address, VLAN, EtherType, and incoming port associated with the packet. This is the default mode.
   - **3 Src/Dest MAC, VLAN, EType, incoming port**: Source and destination MAC addresses, VLAN, EtherType, and incoming port associated with the packet.
   - **4 Src IP and Src TCP/UDP Portfields**: Source IP address and source TCP/UDP port fields of the packet.
   - **5 Dest IP and Dest TCP/UDP Port fields**: Destination IP address and destination TCP/UDP port fields of the packet.
   - **6 Src/Dest IP and TCP/UDP Port fields**: Source and destination IP addresses and source and destination TCP/UDP port fields of the packet.
   - **7 Enhanced hashing mode**: Features modulo-n (mod-n) operation based on the number of ports in the LAG, non-unicast traffic and unicast traffic hashing using a common hash algorithm, excellent load balancing performance, and packet attributes selection based on the packet type:
     - For L2 packets, the source MAC address and destination MAC address are used for hash computation.
     - For L3 packets, the source IP address, destination IP address, and TCP or UDP ports are used for hash computation.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the non-unicast global hash mode

A hash is a fixed-length string of characters that is created from a longer message or data. Hashing is the process of creating a hash, and a hash function is the specific algorithm used to create the hash. Non-unicast global hash mode enhances load balancing for non-unicast traffic, such as multicast and broadcast packets. It allows network devices to distribute non-unicast traffic across multiple links in a LAG using a hash algorithm.

To change the hash mode for non-unicast traffic:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Switching > LAG > LAG Configuration**.
   The LAG Configuration page displays.
3. In the Non Unicast Global Hash Mode section, select the hash mode for non-unicast traffic:
   - **Source Port**: Select the **Source Port** check box to associate the source port number with the packet.
   - **Source MAC/IP Address**: Select the **Source MAC/IP Address** check box to associate the MAC address and IP address of the source port with the packet. By default, this check box is selected.
   - **Destination MAC/IP Address**: Select the **Destination MAC/IP Address** check box to associate the MAC address and IP address of the destination port with the packet. By default, this check box is selected.
   - **Enhanced**: Select the enhanced check box to distribute non-unicast traffic evenly, or load balance the traffic across the LAG. By default, the enhanced mode is disabled.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the settings for a LAG

You can group one or more full-duplex Ethernet links to be aggregated together to form a link aggregation group, which is also known as a port-channel. The switch treats the LAG as if it were a single link.

To configure the settings for a LAG:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > LAG > LAG Configuration**.
   The LAG Configuration page displays.
3. Select the check box for a LAG.
4. To change the default name of the LAG, in the **LAG Name** field, enter the name to be assigned to the LAG.
   By default, the names are ch1, ch2, ch3, and so on. You can enter a name of up to 15 characters.
   For information about adding a name in the Description field, see Add port, LAG, and VLAN descriptions on page 293 or Configure a single LAG and its membership on page 305.

5. From the **Admin Mode** menu, select to enable or disable the LAG.
   When the LAG is disabled, no traffic flows and LACPDUs are dropped, but the links that form the LAG are not released. The default is Enable.

6. From the **Hash Mode** menu, select the hash mode (that is, the load-balancing mode) for the LAG:
   - **1 Src MAC, VLAN, EType, incoming port**: Source MAC address, VLAN, EtherType, and incoming port associated with the packet.
   - **2 Dest MAC, VLAN, EType, incoming port**: Destination MAC address, VLAN, EtherType, and incoming port associated with the packet. This is the default mode.
   - **3 Src/Dest MAC, VLAN, EType, incoming port**: Source and destination MAC addresses, VLAN, EtherType, and incoming port associated with the packet.
   - **4 Src IP and Src TCP/UDP Port fields**: Source IP address and source TCP/UDP port fields of the packet.
   - **5 Dest IP and Dest TCP/UDP Port fields**: Destination IP address and destination TCP/UDP port fields of the packet.
   - **6 Src/Dest IP and TCP/UDP Port fields**: Source and destination IP addresses and source and destination TCP/UDP port fields of the packet.
   - **7 Enhanced hashing mode**: Features MODULO-N operation based on the number of ports in the LAG, non-unicast traffic and unicast traffic hashing using a common hash algorithm, excellent load balancing performance, and packet attributes selection based on the packet type
     - For L2 packets, the source MAC address and destination MAC address are used for hash computation.
     - For L3 packets, the source IP address, destination IP address, and TCP or UDP ports are used for hash computation.

   NOTE: The switch balances traffic on a LAG by selecting one of the links in the channel over which packets must be transmitted. The switch selects the link by creating a binary pattern from selected fields in a packet and associating that pattern with a particular link. The hash mode determines which fields in a packet the switch selects.

7. From the **STP Mode** menu, select to enable or disable the Spanning Tree Protocol (STP) administrative mode for the LAG.
   The default is Enable.

8. From the **Static Mode** menu, select to enable or disable the LAG as a static LAG.
   When the static mode is enabled, the LAG does not transmit or process incoming LACPDUs. That is, a LAG member port does not transmit LACPDUs, and LACPDUs that it receives are dropped. The default is Disable.

9. From the **Link Trap** menu, select to enable or disable the transmission of a trap when the LAG link status changes.
   The default is Enable, which allows a trap to be sent.
10. From the **Local Preference Mode** menu, select if known unicast traffic that is sent to the LAG uses only the LAG interface on the local unit in a stacking environment:
    - **Enable**: This mode is intended for a stacking environment. If the LAG is formed with ports from across the stacking units, any *known* unicast traffic that is sent to the LAG uses only the LAG interface on the local unit. Enable this mode to ensure that the known unicast traffic that is sent to the LAG does not cross the external stack link when the LAG has one or more members on the local unit. This mode does not affect *unknown* unicast, broadcast, and multicast traffic.
    - **Disable**: Known unicast traffic that is sent to the LAG is not restricted to the LAG interface on the local unit. If you do not use stacking, keep this mode disabled, which is the default.
11. Click the **Apply** button.
    Your settings are saved.
12. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 94.LAG configuration information

Field	Description
LAG Description	The description that you might have added to the LAG (see Add port, LAG, and VLAN descriptions on page 293 or Configure a single LAG and its membership on page 305)
LAG ID	The LAG ID
Configured Ports	The interfaces that are members of the LAG
Down Ports	The interfaces that are members of the LAG but that are down
LAG State	Indicates if the LAG link is up or down
Automatic LAG	Indicates if the LAG is manually (No) or automatically (Yes) created
AVB Profile Mode	Indicates the operational status of the AVB profile mode on the LAG

Configure a single LAG and its membership

You can configure a single LAG and add physical interfaces as members to the LAG.
The information about configuring the LAG is identical to the information in Configure the settings for a LAG on page 302, but the information about adding members to the LAG is unique in the following section.

To configure a single LAG and its membership:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> LAG > LAG Membership**.
   The LAG Membership page displays.
3. From the **LAG ID** to menu, select the LAG.
4. To change the default name of the LAG, in the **LAG Name** field, enter the name to be assigned to the LAG.
   By default, the names are ch1, ch2, ch3, and so on. You can enter a name of up to 15 characters.
5. To enter a description for the LAG, in the **LAG Description** field, enter a description of up to 64 characters.
6. From the **Admin Mode** menu, select to enable or disable the LAG.
   When the LAG is disabled, no traffic flows and LACPDUs are dropped, but the links that form the LAG are not released. The default is Enable.
7. From the **Link Trap** menu, select to enable or disable the transmission of traps when the LAG link status changes.
   The default is Disable, which prevents traps from being sent.
8. From the **STP Mode** menu, select to enable or disable the Spanning Tree Protocol (STP) administrative mode for the LAG.
   Enable is the default.
9. From the **Static Mode** menu, select to enable or disable the LAG as a static LAG.
   When the static mode is enabled, the LAG does not transmit or process incoming LACPDUs. That is, a LAG member port does not transmit LACPDUs, and LACPDUs that it receives are dropped. The default is Disable.

10. From the **Hash Mode** menu, select the load-balancing mode for the LAG:
    - **Src MAC, VLAN, EType, incoming port**: Source MAC address, VLAN, EtherType, and incoming port associated with the packet.
    - **Dest MAC, VLAN, EType, incoming port**: Destination MAC address, VLAN, EtherType, and incoming port associated with the packet. This is the default mode.
    - **Src/Dest MAC, VLAN, EType, incoming port**: Source and destination MAC addresses, VLAN, EtherType, and incoming port associated with the packet.
    - **Src IP and Src TCP/UDP Port fields**: Source IP address and source TCP/UDP port fields of the packet.
    - **Dest IP and Dest TCP/UDP Port fields**: Destination IP address and destination TCP/UDP port fields of the packet.
    - **Src/Dest IP and TCP/UDP Port fields**: Source and destination IP addresses and source and destination TCP/UDP port fields of the packet.
    - **Enhanced hashing mode**: Features MODULO-N operation based on the number of ports in the LAG, non-unicast traffic and unicast traffic hashing using a common hash algorithm, excellent load balancing performance, and packet attributes selection based on the packet type
        - For L2 packets, the source MAC address and destination MAC address are used for hash computation.
        - For L3 packets, the source IP address, destination IP address, and TCP or UDP ports are used for hash computation.

> **NOTE:** The switch balances traffic on a LAG by selecting one of the links in the channel over which packets must be transmitted. The switch selects the link by creating a binary pattern from selected fields in a packet and associating that pattern with a particular link.

11. In the Ports table, click each port that you want to make a member of the LAG.
    If stacking is enabled, a Ports table displays for each switch in the stack.
12. Click the **Apply** button.
    Your settings are saved.
13. To save the settings to the running configuration, click the **Save** icon.

MLAG overview

A multi-chassis link aggregation (MLAG), also called Virtual Port Channel (VPC), allows multiple switches to collaborate as a unified logical switch. An MLAG allows links that are physically connected to two different devices to appear as a single port channel to a downstream device. The downstream device connects through a standard LACP link aggregation whose member

ports are split across the two MLAG peer switches.
A VPC domain refers to a pair of peer switches, using the same domain ID (1–255). A peer-link is a dedicated LAG that connects the two peers directly. It carries MLAG control traffic and forwards data only during failover.

NOTE: The MLAG feature is supported on the following M4350 models only:

- M4350-24F4X
- M4350-24F4V
- M4350-24X8F8V
- M4350-32F8V
- M4350-16V4C
- M4350-40F4C
- M4350-16C

You can configure a port channel on a peer link to carry inter-switch control and data traffic between the switches. The port channel partner device is not aware that it is connected over a port channel to two peer switches. Instead, the two switches appear as a single switch to the partner with a single MAC address. All links can carry traffic across a physically diverse topology, and if a link or switch fails, traffic continues flowing with minimal disruption.
MLAG peers function in either a primary or secondary role based on the keepalive priority. The primary switch maintains and propagates spanning-tree and link-aggregation state to the secondary switch.

NOTE: You must disable the Auto-LAG feature on the switch before configuring MLAG features. For more information, see Enable or disable Auto-LAGs on page 299.

Before configuring an MLAG on your switch, note the following requirements:

- Use the same model and firmware on both MLAG peer switches.
- Configure the same VPC domain ID on both peers.
- Use the same LAG number and the same VPC identifier on both peers.
- Use a direct connection between the two peers for the peer-link, with at least two physical ports on each side for redundancy.
- Connect end devices either through an MLAG member LAG or outside the MLAG domain. Single-homeddevices connected to only one peer are not supported behind an MLAG topology.
- Make sure that the two management IP addresses can reach each other to allow peer-keepalive to work. Peer-keepalive, also called peer detection, detects peer failure and avoids dual-active scenarios between the peer switches.
- Configure the peer-link and the MLAG members as dynamic (LACP) LAGs. This is the default LAG type.

Configure the basic VPC global settings

Use the basic VPC Global Configuration page to configure an MLAG on the switch, set the MLAG domain ID, configure the shared VPC MAC address, and adjust keepalive and peer detection settings.
Before you can configure the basic VPC global settings, you need to create a LAG on the switch. For more information, see Configure a single LAG and its membership on page 305.

NOTE: You must configure the VPC MAC address, VPC system priority, and domain ID identically on both MLAG peer switches. If these values differ between peers, the MLAG is not operational.

To configure the basic VPC global settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > MLAG > Basic > VPC Global Configuration**.
   The Basic VPC Global Configuration page displays.
3. In the VPC Global Configuration section, select the VPC Mode **Enable** or **Disable** radio button.
   Select the **Enable** radio button to globally activate the MLAG feature on the switch.
   By default, MLAG is disabled.

4. In the VPC Global Configuration section, configure the following settings:
   - **Domain ID**: Enter a domain ID for the MLAG domain in the range of 1 to 255.

     NOTE: To delete the current domain ID, select the **Remove Domain ID** check box and click the **Apply** button.

   - **VPC MAC Mode**: Select the radio button to let the switch automatically generate the VPC MAC address using the domain ID, or select the radio button to specify a MAC address. The default is **Auto**.
     - **Auto**: The switch automatically generates the VPC MAC address using the domain ID.
     - **Manual**: You can specify a MAC address for the VPC domain in the **VPC MAC** field in xx:xx:xx:xx:xx:xx format.
       The MAC address must be a unicast address, must not match the physical MAC address of either peer switch, and must match on both peer switches. The VPC MAC address range is C0:FF:D4:A7:DA:00 to C0:FF:D4:A7:DA:FF.
   - **VPC System Priority**: Enter a priority value for the VPC domain in the range of 1 to 65535. The default is 32767.

The following table describes the view-only fields in the VPC Global Configuration section.

Table 95. VPC global configuration information

Field	Description
Operational VPC Mode	The current operational mode of VPC on the switch.
VPC State	The current state of the MLAG domain. The VPC state can be one of the following: ○ **Disabled**: The VPC state is disabled on this switch, or the switch has not yet established communication with a peer. ○ **Listen**: VPC is enabled and the switch establishes keepalive communication with its peer. This is a transitional state where no role is elected yet. ○ **Ready**: The peer switches can communicate and function as either the primary or secondary role, but the VPC is not yet fully operational. ○ **Primary**: The VPC is operational and this switch functions as the primary switch. ○ **Secondary**: The VPC is operational and this switch functions as the secondary switch.
Operational VPC MAC	The VPC MAC address the MLAG domain currently uses.
Operational VPC System	The VPC system priority the switch currently uses.

Field	Description
Priority	
Self Role	The role of this switch in the MLAG domain, elected after the switch establishes the keepalive link. The role can be one of the following: ○ **Primary**: The switch acts as the primary switch and maintains and propagates spanning-tree and link-aggregation state to the secondary. **NOTE:** The peer switch with the lower keepalive priority value becomes the primary switch. If both switches have equal priority values, the switch with the lower value MAC address becomes the primary switch. ○ **Secondary**: The switch acts as the secondary switch. ○ **None**: There is no elected role on the switch.
Local System MAC	The MAC address of this switch.

5. In the Keepalive Parameters section, configure the following settings:
   - **Keepalive Priority**: Enter the priority value to use for primary and secondary role selection in the range of 1 to 255. The default is 100.
   - **Keepalive Timeout**: Enter the keepalive timeout in the range of 2 to 15 seconds. If the switch does not receive keepalive messages from the peer within this timeout period, the switch begins the process of transitioning its role. The default is 5 seconds.
   - **Keepalive Admin Mode**: Select the **Enable** or **Disable** radio button.
     - **Enable**: Activates the peer keepalive protocol.
     - **Disable**: Deactivates the peer keepalive protocol. This is the default setting.

The following table describes the view-only fields in the Keepalive Parameters section.

Table 96.Keepalive parameters information

Field	Description
Operational Keepalive Timeout	The keepalive timeout the peer keepalive state machine currently uses while keepalive is enabled. This value can differ from the configured keepalive timeout until you click the **Apply** button.
Keepalive Operational Mode	The current operational mode of the keepalive protocol.

6. In the Peer Link section, configure the following settings:
   - **Enable Modification**: Select the **Enable Modification** check box to allow changes to the peer link configuration.
   - **Port Channel**: From the menu, select the LAG interface between 1 and 64 to configure as the VPC peer link. Select **None** to remove the LAG from the peer link.
7. In the Peer Link section, configure the following settings:
   - **Enable Modification**: Select the **Enable Modification** check box to allow changes to the peer link configuration.
   - **Port Channel**: From the menu, select the LAG interface between 1 and 64 to configure as the VPC peer link. Select **None** to remove the LAG from the peer link.

   The following table describes the view-only fields in the Peer Link section.

Table 97.Peer link information

Field	Description
Peer Link Status	The current status of the peer link.
Peer Keepalive Priority	The keepalive priority of the peer switch.
Peer Link STP Mode	The Spanning Tree Protocol administrative mode the peer link LAG currently uses.

8. In the Peer Detection section, configure the following settings:
   - **Peer Detection Mode**: Select the **Enable** or **Disable** radio button.
     - **Enable**: Peer detection enables the switch to detect peer switches and suppresses state transitions out of the secondary state when the peer link fails.
     - **Disable**: This is the default setting.
   - **Peer IP Address**: Enter the IP address of the peer switch in X.X.X.X format.
   - **Source IP Address**: Enter the local source IP address to use for peer detection packets in X.X.X.X format.
   - **UDP Port**: Enter the UDP port number to use to listen for peer detection packets in the range of 1 to 65535. The default is 50000.
   - **Tx Interval**: Enter the peer detection transmission interval in the range of 200 to 4000 milliseconds. The default is 1000 milliseconds.
   - **Rx Timeout**: Enter the peer detection reception timeout in the range of 700 to 14000 milliseconds. The default is 3500 milliseconds.



NOTE: The Rx Timeout value must be greater than three times the Tx Interval value, and must be less than the keepalive timeout value (in milliseconds).

The following table describes the view-only fields in the Peer Detection section.

Table 98.Peer detection information

Field	Description
Peer Detection Status	The current status of peer detection.
Operational Tx Interval	The peer detection transmission interval the switch currently uses.
Operational Rx Timeout	The peer detection reception timeout the switch currently uses.

9. The following table describes the view-only fields in the Peer section.

Table 99.Peer information

Field	Description
Domain ID	The domain ID the peer switch uses.
VPC MAC	The VPC MAC address the peer switch uses.
Operational VPC MAC	The VPC MAC address the peer switch currently uses.
VPC System Priority	The VPC system priority the peer switch uses.
Operational VPC System Priority	The VPC system priority the peer switch currently uses.
Peer Role	The role of the peer switch in the MLAG domain. Possible values are Primary, Secondary, and None.
Local System MAC	The MAC address of the peer switch.

10. Click the **Apply** button.
    Your settings are saved.
11. To save the settings to the running configuration, click the **Save** icon.

Configure the advanced VPC global settings

Use the advanced VPC Global Configuration page to configure an MLAG on the switch, set the MLAG domain ID, configure the shared VPC MAC address, and adjust keepalive and peer detection settings.
Before you can configure the advanced VPC global settings, you need to create a LAG on the switch. For more information, see Configure a single LAG and its membership on page 305.

NOTE: You must configure the VPC MAC address, VPC system priority, and domain ID identically on both MLAG peer switches. If these values differ between peers, the MLAG is not operational.

To configure the advanced VPC global settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > MLAG > Advanced > VPC Global Configuration**.
   The Advanced VPC Global Configuration page displays.

3. In the VPC Global Configuration section, configure the following settings:
   - **Domain ID**: Enter a domain ID for the MLAG domain in the range of 1 to 255.

     NOTE: To delete the current domain ID, select the **Remove Domain ID** check box and click the **Apply** button.

   - **VPC MAC Mode**: Select the radio button to let the switch automatically generate the VPC MAC address using the domain ID, or select the radio button to specify a MAC address. The default is **Auto**.
     - **Auto**: The switch automatically generates the VPC MAC address using the domain ID.
     - **Manual**: You can specify a MAC address for the VPC domain in the **VPC MAC** field in xx:xx:xx:xx:xx:xx format.
       The MAC address must be a unicast address, must not match the physical MAC address of either peer switch, and must match on both peer switches. The VPC MAC address range is C0:FF:D4:A7:DA:00 to C0:FF:D4:A7:DA:FF.
   - **VPC System Priority**: Enter a priority value for the VPC domain in the range of 1 to 65535. The default is 32767.

The following table describes the view-only fields in the VPC Global Configuration section.

Table 100. VPC global configuration information

Field	Description
Operational VPC Mode	The current operational mode of VPC on the switch.
VPC State	The current state of the MLAG domain. The VPC state can be one of the following: ○ **Disabled**: The VPC state is disabled on this switch, or the switch has not yet established communication with a peer. ○ **Listen**: VPC is enabled and the switch establishes keepalive communication with its peer. This is a transitional state where no role is elected yet. ○ **Ready**: The peer switches can communicate and function as either the primary or secondary role, but the VPC is not yet fully operational. ○ **Primary**: The VPC is operational and this switch functions as the primary switch. ○ **Secondary**: The VPC is operational and this switch functions as the secondary switch.
Operational VPC MAC	The VPC MAC address the MLAG domain currently uses.
Operational VPC System	The VPC system priority the switch currently uses.

Field	Description
Priority	
Self Role	The role of this switch in the MLAG domain, elected after the switch establishes the keepalive link. The role can be one of the following: ○ **Primary**: The switch acts as the primary switch and maintains and propagates spanning-tree and link-aggregation state to the secondary. **NOTE:** The peer switch with the lower keepalive priority value becomes the primary switch. If both switches have equal priority values, the switch with the lower value MAC address becomes the primary switch. ○ **Secondary**: The switch acts as the secondary switch. ○ **None**: There is no elected role on the switch.
Local System MAC	The MAC address of this switch.

4. In the Keepalive Parameters section, configure the following settings:
   - **Keepalive Priority**: Enter the priority value to use for primary and secondary role selection in the range of 1 to 255. The default is 100.
   - **Keepalive Timeout**: Enter the keepalive timeout in the range of 2 to 15 seconds. If the switch does not receive keepalive messages from the peer within this timeout period, the switch begins the process of transitioning its role. The default is 5 seconds.
   - **Keepalive Admin Mode**: Select the **Enable** or **Disable** radio button.
     - **Enable**: Activates the peer keepalive protocol.
     - **Disable**: Deactivates the peer keepalive protocol. This is the default setting.

The following table describes the view-only fields in the Keepalive Parameters section.

Table 101.Keepalive parameters information

Field	Description
Operational Keepalive Timeout	The keepalive timeout the peer keepalive state machine currently uses while keepalive is enabled. This value can differ from the configured keepalive timeout until you click the **Apply** button.
Keepalive Operational Mode	The current operational mode of the keepalive protocol.

5. In the Peer Link section, configure the following settings:
   - **Enable Modification**: Select the **Enable Modification** check box to allow changes to the peer link configuration.
   - **Port Channel**: From the menu, select the LAG interface between 1 and 64 to configure as the VPC peer link. Select **None** to remove the LAG from the peer link.
6. In the Peer Link section, configure the following settings:
   - **Enable Modification**: Select the **Enable Modification** check box to allow changes to the peer link configuration.
   - **Port Channel**: From the menu, select the LAG interface between 1 and 64 to configure as the VPC peer link. Select **None** to remove the LAG from the peer link.

   The following table describes the view-only fields in the Peer Link section.

Table 102.Peer link information

Field	Description
Peer Link Status	The current status of the peer link.
Peer Keepalive Priority	The keepalive priority of the peer switch.
Peer Link STP Mode	The Spanning Tree Protocol administrative mode the peer link LAG currently uses.

7. In the Peer Detection section, configure the following settings:
   - **Peer Detection Mode**: Select the **Enable** or **Disable** radio button.
     - **Enable**: Peer detection enables the switch to detect peer switches and suppresses state transitions out of the secondary state when the peer link fails.
     - **Disable**: This is the default setting.
   - **Peer IP Address**: Enter the IP address of the peer switch in X.X.X.X format.
   - **Source IP Address**: Enter the local source IP address to use for peer detection packets in X.X.X.X format.
   - **UDP Port**: Enter the UDP port number to use to listen for peer detection packets in the range of 1 to 65535. The default is 50000.
   - **Tx Interval**: Enter the peer detection transmission interval in the range of 200 to 4000 milliseconds. The default is 1000 milliseconds.
   - **Rx Timeout**: Enter the peer detection reception timeout in the range of 700 to 14000 milliseconds. The default is 3500 milliseconds.



NOTE: The Rx Timeout value must be greater than three times the Tx Interval value, and must be less than the keepalive timeout value (in milliseconds).

The following table describes the view-only fields in the Peer Detection section.

Table 103.Peer detection information

Field	Description
Peer Detection Status	The current status of peer detection.
Operational Tx Interval	The peer detection transmission interval the switch currently uses.
Operational Rx Timeout	The peer detection reception timeout the switch currently uses.

8. The following table describes the view-only fields in the Peer section.

Table 104.Peer information

Field	Description
Domain ID	The domain ID the peer switch uses.
VPC MAC	The VPC MAC address the peer switch uses.
Operational VPC MAC	The VPC MAC address the peer switch currently uses.
VPC System Priority	The VPC system priority the peer switch uses.
Operational VPC System Priority	The VPC system priority the peer switch currently uses.
Peer Role	The role of the peer switch in the MLAG domain. Possible values are Primary, Secondary, and None.
Local System MAC	The MAC address of the peer switch.

9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

Add a VPC interface configuration

Use the VPC Interface Configuration page to bind a LAG interface to a VPC identifier on the switch.

To add a VPC interface configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > MLAG > Advanced > VPC Interface Configuration**.
   The advanced VPC Interface Configuration page displays.
3. From the **LAG Interface** menu, select the LAG interface to configure as a VPC interface.
4. In the **VPC Identifier** field, enter a VPC identifier in the range of 1 to 63.

5. Click the **Add** button.
   Your settings are saved.
   The **State** field displays the status of the configured VPC interfaces. The status of a VPC interface can be one of the following:
   - **Active**: The VPC is operational. The LAG is bound to the same VPC identifier on both peers and member links are working.
   - **Inactive**: The VPC is configured but not operational. For example, member ports are down on one or both peers.
   - **Disabled**: The VPC is not operational at a configuration level, but the binding exists.
6. To save the settings to the running configuration, click the **Save** icon.

Delete a VPC interface configuration

Use the VPC Interface Configuration page to delete a LAG interface as a VPC interface on the switch.

To delete a VPC interface configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > MLAG > Advanced > VPC Interface Configuration**.
   The advanced VPC Interface Configuration page displays.
3. Select the check box next to the LAG interface that you want to delete. If you want to delete all entries, select the checkbox in the header.
4. Click the **Delete** button.

   NOTE: The changes take effect immediately when you click the **Delete** button.

   Your settings are saved and the VPC interface configuration is deleted.
   If you selected the checkbox in the header, all VPC interface configurations are deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Display VPC interface details

To display the VPC interface details:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Switching > MLAG > Advanced > VPC Interface Details**.
   The advanced VPC Interface Details page displays.
3. From the **LAG Interface** menu, select the LAG interface for which you want to display details.
   View the following information:
   - **Status**: The status of the VPC identifier.
   - **State**: The current operational mode of the VPC interface.
   - **VPC Identifier**: The VPC identifier in the range of 1 to 63.
   - **Configured VLANs**: The VLAN ID or list of VLAN IDs of which the LAG is a member.
   - **Egress Tagging**: The VLAN ID or list of VLAN IDs on which the LAG is tagged.
   - **Peer Link STP Mode**: The Spanning Tree Protocol administrative mode the peer link LAG currently uses.

   The Self Members section displays the following view-only information:
   - **Self Port**: The member ports on this switch.
   - **Status**: The status of each self member port.

   The Peer Members section displays the following view-only information:
   - **Peer Port**: The member ports on the peer switch.
   - **Status**: The status of each peer member port.
4. To refresh the page with the latest information from the switch, click the **Refresh** button.

Display VPC keepalive statistics

To display the VPC keepalive statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > MLAG > Advanced > VPC Keepalive Statistics**.
   The advanced VPC Keepalive Statistics page displays.
3. View the following information:
   - **Total Transmitted**: The total number of messages transmitted with peer keepalive.
   - **Tx Successful**: The total number of messages successfully transmitted with peer keepalive.
   - **Tx Errors**: The total number of message transmit errors with peer keepalive.
   - **Total Received**: The total number of messages received with peer keepalive.
   - **Rx Successful**: The total number of messages successfully received with peer keepalive.
   - **Rx Errors**: The total number of message receive errors with peer keepalive.
   - **Timeout Counter**: The total number of timeouts for peer keepalive.
4. To clear all keepalive statistics and reset them to zero, click the **Clear** button.

5. To refresh the page with the latest information from the switch, click the **Refresh** button.

Display VPC peer link statistics

To display the VPC peer link statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > MLAG > Advanced > VPC Peer Link Statistics**.
   The advanced VPC Peer Link Statistics page displays.

3. View the following information:
   - **Peer Link Control Messages Transmitted**: The total number of control messages transmitted on the peer link.
   - **Peer Link Control Messages Tx Errors**: The total number of control message transmit errors on the peer link.
   - **Peer Link Control Messages Tx Timeout**: The total number of control message transmit timeouts on the peer link.
   - **Peer Link Control Messages ACK Transmitted**: The total number of control message ACKs transmitted on the peer link.
   - **Peer Link Control Messages ACK Tx Errors**: The total number of control message ACK transmit errors on the peer link.
   - **Peer Link Control Messages Received**: The total number of control messages received on the peer link.
   - **Peer Link Data Messages Transmitted**: The total number of data messages transmitted on the peer link.
   - **Peer Link Data Messages Tx Errors**: The total number of data message transmit errors on the peer link.
   - **Peer Link Data Messages Tx Timeout**: The total number of data message transmit timeouts on the peer link.
   - **Peer Link Data Messages Received**: The total number of data messages received on the peer link.
   - **Peer Link BPDU's Transmitted To Peer**: The total number of BPDUs transmitted to the peer switch on the peer link.
   - **Peer Link BPDU's Tx Errors**: The total number of BPDU transmit errors on the peer link.
   - **Peer Link BPDU's Received From Peer**: The total number of BPDUs received from the peer switch on the peer link.
   - **Peer Link BPDU's Rx Errors**: The total number of BPDU receive errors on the peer link.
   - **Peer Link LACPDU's Transmitted To Peer**: The total number of LACPDUs transmitted to the peer switch on the peer link.
   - **Peer Link LACPDU's Tx Errors**: The total number of LACPDU transmit errors on the peer link.
   - **Peer Link LACPDU's Received From Peer**: The total number of LACPDUs received from the peer switch on the peer link.
   - **Peer Link LACPDU's Rx Errors**: The total number of LACPDU receive errors on the peer link.
4. To clear all peer link statistics and reset them to zero, click the **Clear** button.
5. To refresh the page with the latest information from the switch, click the **Refresh** button.

802.1AS timing and synchronization

802.1AS timing and synchronization is an audio video bridging (AVB) feature.
The IEEE 802.1AS standard specifies the protocol and procedures used to ensure that the QoS requirements are guaranteed for time-sensitive applications, such as audio and video.
The IEEE 1588 Precision Time Protocol (PTP) forms the basis of the IEEE 802.1AS standard. PTP specifies a precise clock synchronization protocol that relies on time-stamped packets. For information about configuring PTP, see Precision Time Protocol transparent clock on page 75 and Precision Time Protocol boundary clock and grandmaster clock on page 78.
The PTP protocol is applicable to distributed systems that consist of one or more nodes communicating over of communication media. The distribution of synchronous time information occurs in a hierarchical manner with a grandmaster clock at the root of the hierarchy. The grandmaster provides a common and precise time reference for one or more directly-attached devices by periodically exchanging timing information. That is, all devices synchronize their clocks with the grandmaster clock. The devices can, in-turn, act as master devices for further hierarchical layers of other devices.

NOTE: 802.1AS audio video bridging (AVB) is not compatible with stacking. You can configure one or the other.

Configure and view the global 802.1AS settings

You can globally enable 802.1AS on the switch and configure local clock priorities. The 802.1AS feature calculates the time delay between devices on a link and maintains an accurate view of a network clock.

To configure and view the global 802.1AS settings on the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > 802.1AS > Basic > 802.1AS Configuration**.
   The 802.1AS Configuration page displays.
3. To enable the feature, select the 802.1AS status **Enable** radio button.
4. In the **Local Clock Priority1** field, specify the first priority value of the local clock.
   The local clock is the clock for the switch as a time-aware bridge. Enter a number between 0 and 255 seconds. The default is 255 seconds.
5. In the **Local Clock Priority2** field, specify the second priority value of the local clock.
   Enter a number between 0 and 255 seconds. The default is 248 seconds.

6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 105.802.1AS configuration information

Field	Description
GrandMaster Capable	Identifies whether the grandmaster clock is active. The default is False.
Best Clock Identity	The MAC address of the best clock that the switch detected.
Best Clock Priority1	The first priority value of the best clock that the switch detected.
Best Clock Priority2	The second priority value of the best clock that the switch detected.
Steps to Best Clock	The number of links between the best clock and the switch. If the switch itself is the best clock, the value is zero.
Local Clock Identity	The MAC address of the local clock on the switch.
GM Change Count	The number of times the grandmaster clock changed.
Last GM Change Timestamp	The time when the most recent grandmaster clock change occurred.

Configure the 802.1AS interface settings

You can configure 802.1AS settings for individual interfaces.

To configure the 802.1AS interface settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > 802.1AS > Advanced > 802.1AS Port Settings.**
   The 802.1AS Port Settings page displays.

3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Admin Mode** menu, select **Enable** or **Disable**.
   For the 802.1AS settings to take effect on the interface, select **Enable**. The default is Disable.
6. In the **Pdelay Threshold** field, specify the propagation delay threshold for the interface.
   The threshold determines whether the port is capable of participating in the 802.1AS protocol. If the propagation delay on the interface is above the threshold you configure, the interface is not considered capable of participating in the 802.1AS protocol. The peer delay must be less than the threshold value configured on the interface. The range is from 0 to 1,000,000,000 nanoseconds (ns). The default is 2500 ns.
7. In the **Allowed Lost Responses** field, specify the allowed loss response value.
   If the interface does not receive valid responses to PDELAY_REQ messages above the value of the allowed lost responses, a port is considered to not be exchanging peer delay messages with its neighbor. The range is from 0 to 65,535. The default is 3.
8. In the **Initial Sync Interval** field, specify the transmission rate of SYNC messages.
   This value is the logarithm to the base 2 of the mean-time interval between successive SYNC messages sent on this interface. The configured initial interval becomes the current interval only after the port is initialized or re-initialized for 802.1AS operation. The range is from –5 to 5. The default is –3.

9. In the **Initial Pdelay Interval** field, specify the transmission rate of PDELAY_REQ messages.
   This value is the logarithm to the base 2 of the mean time interval between successive PDELAY_REQ messages sent on this interface. The configured initial interval becomes the current interval only after the port is initialized or re-initialized for 802.1AS operation. The range is from –5 to 5. The default is –3.
10. In the **Initial Announce Interval** field, specify the transmission rate of ANNOUNCE messages.
    This value is the logarithm to the base 2 of the mean time interval between successive ANNOUNCE messages sent on this interface. The configured initial interval becomes the current interval only after the port is initialized or re-initialized for 802.1AS operation. The range is from –5 to 5. The default is –3.
11. In the **SyncRx Timeout** field, specify the number of SYNC intervals.
    This value sets the number of SYNC intervals that must pass without receipt of SYNC information before the switch detects that the master clock is no longer transmitting.
12. n the **AnnounceRx Timeout** field, specify the number of ANNOUNCE intervals.
    This value sets the number of ANNOUNCE intervals that must to pass without receipt of ANNOUNCE PDU before the switch detects that the master clock is no longer transmitting.
13. Click the **Apply** button.
    Your settings are saved.
14. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 106.802.1AS interface information

Field	Description
Port Role	The 802.1AS role of the interface. The possible roles are as follows: • Disabled (default) • Master • Slave • Passive
Propagation Delay	The mean propagation delay on the interface.
Measuring Pdelay	Indicates whether the interface is receiving PDELAY response messages from the other end of the link.
802.1AS Capable	Indicates whether the interface is 802.1AS capable. By default, the interface is not 802.1AS capable.
Neighbor Rate Ratio	An estimated ratio of the frequency of the local clock entity of the time-aware device at the other end of the link of the interface in relation to the frequency of the local clock entity of the switch.
Current Sync Interval	The current mean time interval between successive SYNC messages sent over the link, in logarithm to base 2 format.
Current Pdelay Interval	The current mean time interval between successive PDELAY_REQ messages sent over the link, in logarithm to base 2 format.
Current Announce Interval	The current mean time interval between successive ANNOUNCE messages sent over the link, in logarithm to base 2 format.

View the 802.1AS statistics

You can view the 802.1AS statistics for the interfaces.

To view the 802.1AS statistics for the interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > 802.1AS > Advanced > 802.1AS Statistics**.

3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1**: Only physical interfaces are displayed. This is the default setting.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed.

   The following table describes the view-only fields on the page.

Table 107.802.1AS statistics information

Field	Description
Interface	The interface for which information is displayed.
Sync Tx	The total number of SYNC packets transmitted without errors.
Sync Rx	The total number of SYNC packets received without errors.
Followup Tx	The total number of FOLLOWUP packets transmitted without errors.
Followup Rx	The total number of FOLLOWUP packets received without errors.
Announce Tx	The total number of ANNOUNCE packets transmitted without errors.
Announce Rx	The total number of ANNOUNCE packets received without errors.
Pdelay Req Tx	The total number of PDELAY_REQ packets transmitted without errors.
Pdelay Req Rx	The total number of PDELAY_REQ packets received without errors.
Pdelay Resp Tx	The total number of PDELAY_RESP packets transmitted without errors.
Pdelay Resp Rx	The total number of PDELAY_RESP packets received without errors.
Pdelay Resp Followup Tx	The total number of PDELAY_RESP_FOLLOWUP packets transmitted without errors.
Pdelay Resp Followup Rx	The total number of PDELAY_RESP_FOLLOWUP packets received without errors.
Signaling Tx	The total number of SIGNALING packets transmitted without errors.
Signaling Rx	The total number of SIGNALING packets received without errors.
Sync Timeouts	The total number of SYNC receipt time-outs that occurred.
Sync Discards	The total number of SYNC packets that were discarded.
Announce Timeouts	The total number of ANNOUNCE receipt time-outs that occurred.
Announce Discards	The total number of ANNOUNCE packets that were discarded.
Pdelay Timeouts	The total number of PDELAY receipt time-outs that occurred.
Pdelay Discards	The total number of PDELAY packets discarded.

Field	Description
Bad Headers	The total number of packets received with a bad header.

Multiple Registration Protocol and 802.1Qav

Multiple Registration Protocol (MRP) is an audio video bridging (AVB) feature. MVR is a base registration protocol that enables devices running an MRP application to register attributes to other devices in a network. MRP provides an application to register attributes such as bandwidth for an audio-video (AV) stream and MAC address information. This feature is used by various applications to propagate the registration.

NOTE: MRP framework must be available and enabled on all intermediate devices to ensure that the propagation of the attributes occurs throughout the network.

The switch supports the following MRP applications:

- **Multiple MAC Registration Protocol (MMRP)**: MMRP allows for the propagation MAC address information in the network, and allows for the registration and deregistration of both individual MAC address information and group MAC address membership. End stations can request to join or leave a multicast group, or to register an individual MAC address with a specific VLAN. MAC address entries can be dynamically registered and deregistered if MMRP is administratively enabled on the switch.
- **Multiple VLAN Registration Protocol (MVRP)**: MVRP registers VLANs in the network, enabling automatic VLAN configuration on the switch. In a typical network, VLAN tagging is common. Many nodes require ingress traffic to be tagged with a specific VLAN ID, and other nodes require egress traffic to be transmitted with a specific VLAN ID. With the use of MVRP on both ingress and egress, no manual VLAN configuration is required to pass tagged traffic through the network.
- **Multiple Stream Reservation Protocol (MSRP)**: MSRP reserves resources in the network to facilitate time-sensitive traffic to flow end to end. A typical network includes multiple talkers (devices that transmit streams) and multiple listeners (devices that receive streams from one or many talkers). Each flow has specific bandwidth, frame rate, and time sync requirements. MSRP guarantees these resources through all intermediate devices between any talker and listener.

With MRP, network attributes are declared, registered, withdrawn, and removed dynamically without user intervention. This dynamic nature is especially useful in networks where the following is true:

- Network attributes are likely to change frequently, requiring reconfiguration of the intermediate devices.
- Recipients of these attributes frequently increase or decrease in number.

- Each of these changes without a dynamic self-adjusting framework would require constant attention from a network administrator.

Configure the global MRP settings

You can configure global MRP settings for the switch.

To configure the global MRP settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > MRP > Basic > MRP Configuration**.
   The MRP Configuration page displays.
3. Select the MVRP Mode **Enable** or **Disable** radio button to specify the global administrative mode of MVRP on the switch.
   Enabling this mode lets MVRP register VLANs in the network, enabling automatic VLAN configuration on the device. The default is Disable.
4. Select the MMRP **Enable** or **Disable** radio button to specify the global administrative mode of MMRP on the switch.
   Enabling this mode allows for the propagation of MAC address information in the network. In turn, this allows for the registration and deregistration of both individual MAC address information and group MAC address membership. The default is Disable.
5. Select the MSRP **Enable** or **Disable** radio button to specify the global administrative mode of MSRP on the switch.
   Enabling this mode allows the switch to reserve resources in the network to facilitate time-sensitive traffic to flow end to end. The default is Disable.
6. Select the MSRP Talker Pruning **Enable** or **Disable** radio button.
   Enabling this mode limits the switch to transmitting talker protocol data units (PDUs) to destination MAC addresses that are in the Multicast Forwarding Table (MFDB). A talker PDU is the source of an audio-video stream. The default is Disable.
7. Select the Periodic State Machine (MVRP Mode) **Enable** or **Disable** radio button.
   Enabling this mode can limit the effect of MVRP topology changes and reduce the number of PDUs transmitted between devices. The default is Disable.
8. Select the Periodic State Machine (MMRP) **Enable** or **Disable** radio button.
   Enabling this mode can limit the effect of MMRP topology changes and reduce the number of protocol data units (PDUs) transmitted between devices. The default is Disable.

9. In the **MSRP Max Fan In Ports** field, specify the maximum number of interfaces on which MSRP registrations are allowed.
   The range depends on the switch model. The default is 0, which means that no limit is imposed on the number of interfaces.
10. Select the MSRP Boundary Propagation **Enable** or **Disable** radio button.
    When you enable MSRP boundary propagation, depending on the configuration, a stream reservation can be allowed to cross a boundary port.
    The default is Disable.
11. In the **MSRP PDU Transmit Time Gap** field, specify the period between the transfer of two consecutive MSRP PDUs after the LeaveAll time expires.
    After the LeaveAll time expires, all participants are deregistered and must rejoin by responding to the PDUs that the switch sends. By sending consecutive MSRP PDUs at an interval, the packet buffer of a receiving endpoint does not overflow. The period between consecutive MSRP PDUs is a value in milliseconds.
12. Click the **Apply** button.
    Your settings are saved.
13. To save the settings to the running configuration, click the **Save** icon.

Configure 802.1Qav mapping

You can set the global IEEE 802.1Qav settings, which are QoS priorities for class A and class B traffic Ethernet audio video (EAV) streams.
The IEEE 802.1Qav standard supports time-sensitive traffic streams by pacing all switch traffic, including legacy asynchronous Ethernet traffic, through queuing and forwarding. When a talker declares a stream, 802.1Qav identifies whether the stream is class A or class B and specifies the stream's bandwidth requirements. Class A traffic receives a higher transmission priority than class B traffic.

To change the global 802.1Qav mapping priorities:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > MRP > Basic > MRP Configuration**.
   The MRP Configuration page displays.
3. Scroll to the 802.1qav Mapping section.
   By default, the EAV priority for class A traffic is 3, with a remap priority of 0.
   By default, the EAV priority for class B traffic is 2, with a remap priority of 0.

4. In the EAV Stream column, select the check box for the class.
5. In the **EAV Priority** field, specify a priority from 0 to 7.
   Priority 0 is the lowest priority; Priority 7 is the highest priority.
6. In the **EAV Remap Priority** field, specify a priority from 0 to 7.
   Priority 0 is the lowest priority; Priority 7 is the highest priority.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Configure the MRP interface settings

You can configure the MRP mode and timer settings for one or more interfaces. The timers control when and how often various messages are transmitted on each interface.

To configure MRP interface settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> MRP > Advanced > MRP Port Settings**.
   The MRP Port Settings page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

5. From the **MVRP Mode** menu, select **Enable** or **Disable** to specify the administrative mode of MVRP on the interface.
   Enabling this mode lets MVRP register VLANs in the network, enabling automatic VLAN configuration on the interface. The default is Disable.
6. From the **MMRP Mode** menu, select **Enable** or **Disable** to specify the administrative mode of MMRP on the interface.
   Enabling this mode allows for the propagation of MAC address information on the interface. In turn, this allows for the registration and deregistration of both individual MAC address information and group MAC address membership. The default is Disable.
7. From the **MSRP Mode** menu, select **Enable** or **Disable** to specify the administrative mode of MSRP on the interface.
   Enabling this mode allows the interface to reserve resources in the network to facilitate time-sensitive traffic to flow end to end. The default is Disable.
8. In the **Join Timer** field, specify the time that the interface must wait for JoinIn messages from other MRV participants after the interface sends a Join message.
   If the time that you specify passes before the interface receives a JoinIn message, the interface resends the Join message. The range is from 10 to 100 centiseconds. The default is 20.
9. In the **Leave Timer** field, specify the time that the interface must wait before it deregisters attributes from other MRV participants.
   If the interface receives Join messages from other participants before the time that you specify expires, the attributes are not deregistered. The range is from 20 to 600 centiseconds. The default is 100.
10. In the **Leave All Timer** field, specify the time that the interface must wait after it starts the MRP registration process before the participants refresh and reregister their attributes. The range is 200 to 6000 centiseconds. The default value is 1000.
11. In the **MSRP SR class PVID** field, specify the default VLAN ID that the interface must use for MSRP stream traffic.
12. Click the **Apply** button.
    Your settings are saved.
13. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 108.MRP interface settings information

Field	Description
MSRP Class A Boundary Port	Indicates whether the interface is a boundary port for class A traffic.
MSRP Class B Boundary Port	Indicates whether the interface is a boundary port for class B traffic.

Display or clear MMRP statistics

You can display or clear information regarding the MMRP frames transmitted and received by the switch and by each interface.

To display or clear MMRP statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > MRP > Advanced > MMRP Statistics**.
   The MMRP Statistics page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. To refresh the page, click the **Refresh** button.
5. To clear the statistics, do the following:
   a. Select the check box next to the interface or interfaces, or, to clear the statistics for all interfaces, select the check box in the table heading.
   b. Click the **Clear** button.
   The statistics are cleared.
6. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 109.MMRP statistics information

Field	Description
MMRP Global Statistics	
Frames Received	The number of MMRP frames that were received on the switch.
Bad Header	The number of MMRP frames with bad headers that were received on the switch.
Bad Format	The number of MMRP frames with bad PDUs body formats that were received on the switch.
Frames Transmitted	The number of MMRP frames that were transmitted on the switch.
Transmission Failures	The number of MMRP frames that the switch did not transmit.
MMRP Statistics	
Interface	The interface for which information is displayed.
Frames Received	The number of MMRP frames that were received on the interface.
Bad Header	The number of MMRP frames with bad headers that were received on the interface.
Bad Format	The number of MMRP frames with bad PDUs body formats that were received on the interface.
Frames Transmitted	The number of MMRP frames that were transmitted on the interface.
Transmission Failures	The number of MMRP frames that the interface failed to transmit.

Display or clear MVRP statistics

You can display or clear information about the MVRP frames transmitted and received by the switch and by each interface.

To display or clear MVRP statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > MRP > Advanced > MVRP Statistics**.
   The MVRP Statistics page displays.

3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. To refresh the page, click the **Refresh** button.
5. To clear the statistics, do the following:
   a. Select the check box next to the interface or interfaces, or, to clear the statistics for all interfaces, select the check box in the table heading.
   b. Click the **Clear** button.
   The statistics are cleared.
6. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 110.MVRP statistics information

Field	Description
MVRP Global Statistics	
Frames Received	The number of MMRP frames that were received on the switch.
Bad Header	The number of MMRP frames with bad headers that were received on the switch.
Bad Format	The number of MMRP frames with bad PDUs body formats that were received on the switch.
Frames Transmitted	The number of MMRP frames that were transmitted on the switch.
Transmission Failures	The number of MMRP frames that the switch did not transmit.
Message Queue Failures	The number of MMRP messages that were not added to the queue.
MVRP Statistics	
Interface	The interface for which information is displayed.
Frames Received	The number of MVRP frames that were received on the interface.
Bad Header	The number of MVRP frames with bad headers that were received on the interface.
Bad Format	The number of MVRP frames with bad PDUs body formats that were received on the interface.
Frames Transmitted	The number of MVRP frames that were transmitted on the interface.
Transmission Failures	The number of MVRP frames that the interface failed to transmit.
Registration Failures	The number of MVRP frames that failed to register on a device or particular interface.

Display or clear MSRP statistics

You can display or clear information about the MSRP frames transmitted and received by the switch and by each interface.

To display or clear MSRP statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Switching > MRP > Advanced > MSRP Statistics**.
   The MSRP Statistics page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. To refresh the page, click the **Refresh** button.
5. To clear the statistics, do the following:
   a. Select the check box next to the interface or interfaces, or, to clear the statistics for all interfaces, select the check box in the table heading.
   b. Click the **Clear** button.
      The statistics are cleared.
6. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 111.MSRP statistics information

Field	Description
MSRP Global Statistics	
Frames Received	The number of MSRP frames that were received on the switch.
Bad Header	The number of MSRP frames with bad headers that were received on the switch.
Bad Format	The number of MSRP frames with bad PDUs body formats that were received on the switch.
Frames Transmitted	The number of MSRP frames that were transmitted on the switch.
Transmission Failures	The number of MSRP frames that the switch did not transmit.
Message Failures	The number of MSRP messages that were not added to the queue.
MSRP Statistics	
Interface	The interface for which information is displayed.
Frames Received	The number of MSRP frames that were received the interface.
Bad Header	The number of MSRP frames with bad header that were received on the interface.
Bad Format	The number of MSRP frames with bad PDUs body format that were received on the interface.
Frames Transmitted	The number of MSRP frames that transmitted on the interface.
Transmission Failures	The number of MSRP frames that the interface failed to transmit.
Registration Failures	The number of MSRP frames that failed to register on a device or particular interface.

Display the MSRP reservation settings

You can display information about the talker, listener, and intermediate device status for the devices involved in each MSRP stream that flows through the switch.

To display the MSRP reservation settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Switching > MRP > Advanced > MSRP Reservation Parameters**.
   The MSRP Reservation Parameters page displays.

3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 112.MSRP reservation parameters information

Field	Description
Interface	The interface for which information is displayed.
Stream ID	A 16-bit unsigned integer value, which is a unique ID used to distinguish the stream among multiple streams from the same device.
Listener Declaration Status	The MSRP declaration status of the listener attribute.
Listener Declaration Type	The MSRP declaration type of the listener attribute.
Talker Declaration Status	The MSRP declaration status of the talker attribute.
Talker Declaration Type	The MSRP declaration type of the talker attribute.
Accumulated Latency	The latency, in nanoseconds, that occurred for the stream on its path from the talker to a listener. The latency starts with zero and is increased by one for each device that the stream propagates through in the network.
Failure Bridge Interface	If a failure occurred, the interface on the device on which the failure occurred.
Failure Bridge MAC	If a failure occurred, the MAC address of the device on which the failure occurred.
Failure Code	If a failure occurred, the number that represents the reason for the failure. The switch supports the following codes: ○ 1: Insufficient bandwidth ○ 3: Insufficient bandwidth for the traffic class ○ 5: Stream destination address is already in use ○ 7: Reported latency changed ○ 8: Egress port is not capable of audio video bridging (AVB) ○ 9: Use a different destination address (for example, if the MAC DA hash table is full) ○ 12: Cannot store destination address (for example, if the device ran out of MAC DA resources) ○ 13: Requested priority is not an SR class priority ○ 14: MaxFrameSize is too large for media ○ 15: msrpMaxFanInPorts limit was reached ○ 16: Changes in FirstValue for a registered StreamID ○ 17: VLAN is blocked on this egress port (Registration Forbidden)

Field	Description
Stream Age	The time, in seconds, since the stream destination address was added to the Dynamic Reservations Entries table. A value of zero indicates the destination address was not added to the table.

Configure and display the Qav settings for interfaces

You can configure and view the IEEE 802.1Qav settings for interfaces. The IEEE 802.1Qav standard supports time-sensitive traffic streams by pacing all switch traffic, including legacy asynchronous Ethernet traffic, through queuing and forwarding. When a talker declares a stream, 802.1Qav identifies whether the stream is class A or class B and specifies the stream's bandwidth requirements. Class a traffic receives a higher transmission priority than class B traffic.
You can configure and display the selected bandwidth allocations for class A and class B traffic.

To configure and display the Qav settings for interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching>MRP > Advanced > Qav Parameters**.
   The Qav Parameters page displays.
3. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
4. In the Class A **MSRP Delta Bandwidth** field, specify the additional bandwidth represented as a percentage of the interface transmit rate that is reserved for class A and class B traffic Class A traffic receives a higher priority than class B traffic. The range is from 0 to 100. The default is 100.
5. In the Class B **MSRP Delta Bandwidth** field, specify the additional bandwidth represented as a percentage of the interface transmit rate that is reserved for class B traffic
   The range is from 0 to 100. The default is 0.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 113.Qav settings information

Field	Description
Class A Bandwidth Allocated	The current rate (in Bps) of class A traffic on interface
Class A Remaining Bandwidth	The maximum rate (in Bps) of class A that is supported by the interface.
Class B Bandwidth Allocated	The current rate (in Bps) of class B traffic on interface
CLass B Remaining Bandwidth	The maximum rate (in Bps) of class B that is supported by the interface.
Total Bandwidth Allocated	The sum of the allocated class A and class B traffic rates (in Bps) on interface.
Total Remaining Bandwidth	75 percent of the interface speed minus that total allocated bandwidth (in Bps/sec).

Display MSRP streams information

You can display information about MSRP streams that are flowing through each interface.

To display MSRP streams information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching > MRP > Advanced > MSRP Streams Information**.
   The MSRP Streams Information page displays.

3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 114.MSRP streams information

Field	Description
Stream ID	A 16-bit unsigned integer value, which is a unique ID used to distinguish the stream among multiple streams from the same device
Stream Source MAC Address	The MAC address of the traffic stream's source
Received Accumulated Latency	The 32-bit unsigned accumulated latency value that is used to determine the worst-case latency that a stream can suffer in its path from the talker to a listener. The latency starts with zero and is increased by one for each device that the stream propagates through in the network.
Traffic Class	Shows whether the stream is class A or class B. Class A traffic receives higher priority than class B traffic.
Rank	The 5-bit unsigned rank value that is used by devices to determine which streams can still be served when the MSRP registrations exceed the capacity of an interface to carry the corresponding data streams. If a device becomes oversubscribed (for example, a network reconfiguration or 802.11 bandwidth reduction occurs), the rank value is also used to determine which stream are dropped. A lower numeric value is more important than a higher numeric value.
TSpec Max Frame Size	The 32-bit unsigned bandwidth value that is used to allocate resources and adjust queue selection settings to supply the quality of service that is requested by an MSRP talker. The value represents the maximum rate, in units of 1024 octets per second, at which frames in the stream referenced by the talker can be transmitted.
TSpec Max Interval Frames	The 32-bit unsigned frame rate value that is used to allocate resources and adjust queue selection settings to supply the quality of service that is requested by an MSRP talker. The value represents the maximum number of frames that the talker can transmit in one second.
Stream VLAN	The VLAN ID of the traffic stream
Destination MAC	The MAC address of the traffic stream's destination
Received Failure Information Bridge Interface	If a failure occurred, the interface of the device on which the failure occurred

Field	Description
Received Failure Information Failure Code	If a failure occurred, the number that represents the reason for the failure. The switch supports the following codes: ○ 1: Insufficient bandwidth ○ 3: Insufficient bandwidth for the traffic class ○ 5: Stream destination address is already in use ○ 7: Reported latency changed ○ 8: Egress port is not capable of audio video bridging (AVB) ○ 9: Use a different destination address (for example, if the MAC DA hash table is full) ○ 12: Cannot store destination address (for example, if the device ran out of MAC DA resources) ○ 13: Requested priority is not an SR class priority ○ 14: MaxFrameSize is too large for media ○ 15: msrpMaxFanInPorts limit was reached ○ 16: Changes in FirstValue for a registered StreamID ○ 17: VLAN is blocked on this egress port (Registration Forbidden)
Received Failure Information Bridge MAC	If a failure occurred, the MAC address of the device on which the failure occurred
Talker Interface	The interface on which the talker is present
Listeners	The interface on which listeners are present

Loop protection

Loop protection can detect physical and logical loops between Ethernet ports on the switch. Loops inside a network can be costly because they consume resources and reduce the performance of the network. Detecting loops manually can be cumbersome. The switch can automatically identify loops in the network. You can enable loop protection per globally and per port.

If loop protection is enabled, the switch sends predefined protocol data unit (PDU) packets to a Layer 2 broadcast destination address (FF:FF:FF:FF:FF:FF) on all ports for which the feature is enabled. You can selectively disable PDU packet transmission for loop protection on specific ports even while port loop protection is enabled. If the switch receives a packet with the previously mentioned broadcast destination address, the source MAC address in the packet is compared with the MAC address of the switch. If the MAC address does not match, the packet is forwarded to all ports that are members of the same VLAN, just like any other broadcast packet. The packet is not forwarded to the port from which it was received

If the source MAC address matches the MAC address of the switch, the switch can perform one of the following actions, depending on how you configure the action:

- The port is shut down.
- A log message is generated. (If a syslog server is configured, the log message can be sent to the syslog server.)
- The port is shut down and a log message is generated.

Loop protection is not intended for ports that serve as uplinks between spanning tree–aware switches. It is intended for unmanaged switches that drop spanning tree BPDUs. Loop protection detects physical and logical loops between Ethernet ports on a device. You must enable loop protection globally before you can enable and configure it at the interface level. Loop protection is supported on physical interfaces and static LAG interfaces, but not on dynamic LAG interfaces.

About loop protection

Loops inside a network are costly because they consume resources and reduce the performance of the network. Detecting loops manually can be cumbersome.

The switch can automatically identify loops in the network. You can enable loop protection per port or globally.

If loop protection is enabled, the switch sends predefined protocol data unit (PDU) packets to a Layer 2 multicast destination address (09:00:09:09:13:A6) on all ports for which the feature is enabled. You can selectively disable PDU packet transmission for loop protection on specific ports even while port loop protection is enabled. If the switch receives a packet with the previously mentioned multicast destination address, the source MAC address in the packet is compared with the MAC address of the switch. If the MAC address does not match, the packet is forwarded to all ports that are members of the same VLAN, just like any other multicast packet. The packet is not forwarded to the port from which it was received.

If the source MAC address matches the MAC address of the switch, the switch can perform one of the following actions, depending on how you configure the action:

- The port is shut down.
- A log message is generated. (If a syslog server is configured, the log message can be sent to the syslog server.)
- The port is shut down and a log message is generated.

If loop protection is disabled, the multicast packet is silently dropped.

Loop protection is not intended for ports that serve as uplinks between spanning tree–aware switches. Loop protection is designed for unmanaged switches that drop spanning tree bridge protocol data units (BPDUs).

You need to enable the feature globally before you can enable it at the port level so that the system policy filter can be installed.

Loop protection treats PDU packets transmission and spanning tree protocol in the following ways:

- **Loop protection and PDU packet transmission**: Loop protection sends loop protocol packets from all ports on which it is enabled. You can configure the interval (1 to 5 seconds)

between two successive loop protection PDU packets. The default interval is 5 seconds. If the switch receives a loop protocol packet on a port for which the action is set to shut down the port, the port can no longer receive and send frames. Loop protection operates at a port level, regardless of VLAN assignment and membership, detecting loops across VLANs.

- **Loop protection and Spanning Tree Protocol**: Loop protection does not impact end nodes and is not intended for ports that serve as uplinks between spanning tree–aware switches. Loop protection can coexist with Spanning Tree Protocol (STP). You can enable both loop protection and STP on a port because these features function independently of each other. STP does not bring a port down when a loop is detected but keeps the port in blocking state. Because PDUs are allowed in a blocking state, loop protection packets are received and loop protection brings down the port that is involved in the loop (if the configured action is to shut down the port).

Configure the global loop protection settings

Before you can configure loop protection for individual ports (see Configure the loop protection settings for interfaces and display the loop protection state on page 348), you must globally enable and configure loop protection.

To globally enable and configure loop protection:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> L2 Loop Protection > L2 Loop Protection Configuration**.
   The L2 Loop Protection Configuration page displays.
3. In the Global L2 Loop Protection Configuration section, configure the following settings:
   - Next to Admin Mode, select the **Enable** or **Disable** radio button to specify the administrative mode of loop protection on the switch. By default, loop protection is globally disabled.
   - From the **TLV Advertized Interval** menu, select the interval between the transmissions of loop packets on a port. The range is from 1 to 5 seconds. The default setting is 5 seconds. The selected interval applies to all ports for which you enable loop protection.
   - From the **Max PDU Receive** menu, select the maximum number of packets that a port can receive before an action is taken. The default setting is 1 packet. The selected number of packets applies to all ports for which you enable loop protection.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the loop protection settings for interfaces and display the loop protection state

Before you can configure loop protection for individual ports, you must globally enable loop protection (see Configure the global loop protection settings on page 347).

To enable and configure loop protection for an interface and display the loop protection state on the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Switching> L2 Loop Protection > L2 Loop Protection Configuration**.
   The L2 Loop Protection Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Keep Alive** menu, select **Enable** to specify that loop protection must be enabled on the port.
   By default, loop protection is disabled for a port.
6. From the **RX Action** menu, select the action that the switch takes when a loop is detected on the port:
   - **Log**: Log the message when a loop is detected on the port.
   - **Disable**: Disable the port when a loop is detected. This is the default setting.
   - **Both**: Log and disable the port when a loop is detected.

7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.
9. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 115.Loop protection interface configuration information

Field	Description
Loop Detected	Indicates (Yes or No) whether a loop is detected on the interface
Loop Count	The number of packets that were received on the interface after the loop was detected
Time Since Last Loop	The time since the loop was detected
Port Status	The status of the interface (Enabled or Disabled)
Errdisable Reason	The reason that the interface was disabled. In addition to being disabled because of loop protection, the interface can be disabled because of Unidirectional Link Detection (UDLD), a broadcast storm, a unicast storm, and so on.
Auto Recovery Time Left (sec)	The time that is left before the interface is reenabled through the autorecovery process. The time is in the range from 30 to 604800 seconds.

Manage Routing

This chapter covers the following topics:

- Routing concepts
- Routing table, routes and route preferences
- IPv4 routing
- IPv6 routing
- Routing VLANs
- Address Resolution Protocol
- Routing Information Protocol
- Router discovery and router advertisements
- Virtual Router Redundancy Protocol

Routing concepts

The switch supports IP routing. When a packet enters the switch, the switch checks the destination MAC address to determine if it matches any of the configured routing interfaces. If it does, the switch searches the host table for a matching destination IP address. If a matching entry is found, the packet is routed to the host. If no matching entry is found, the switch performs a longest prefix match on the destination IP address. If a matching entry is found, the packet is routed to the next hop. If no matching entry is found, the packet is routed to the next hop that is specified in the default route. If no default route exists, the packet is dropped.

Routing table, routes and route preferences

The routing table collects routes from multiple sources and list all routes: static routes, local routes, dynamically added routes, and so on. Static routes are routes that you manually add. The routing table can learn multiple routes to the same destination from multiple sources, for example dynamically added routes through routing protocols.

Configure a route and display learned routes

You can add or change a route in the routing table and display the automatically learned routes in the routing table.

To add or change a route and display learned routes:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Routing > Routing Table > Basic > Route Configuration**.
   The Route Configuration page displays.
3. If you are changing an existing route, select the check box for the route.
4. From the **Route Type** menu, select one of the following route types:
   - **Default**: Specify the next hop address and preference.
   - **Static**: Specify the network address, subnet mask, next hop address, and preference.
   - **Static Reject**: Specify the network address, subnet mask, and preference.
5. In the **Network Address** field, enter the IP route prefix for the destination.
6. In the **Subnet Mask** field, enter the portion of the IP interface address that identifies the attached network.
   This is also referred to as the subnet/network mask.
7. In the **Next Hop IP Address**, enter the IP address of the outgoing router that must be used when traffic is forwarded to the next router (if any) in the path toward the destination.
   The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly attached network.
8. In the **Preference** field, enter a value from 1 to 255.
   You can specify the preference value (sometimes called administrative distance) of an individual static route. Among routes to the same destination, the route with the lowest preference value is the route entered into the forwarding database. By specifying the preference of a static route, you can control if a static route is more preferred or less preferred than routes from dynamic routing protocols. The preference also controls if a static route is more preferred or less preferred than other static routes to the same destination.
9. In the **Description** field, enter the description of this route that identifies the route.
   The description must consist of alphanumeric, hyphen, or underscore characters and can be up to 31 characters in length.
10. Do one of the following:
    - If you are adding a new route, click the **Add** button. Your settings are saved. The route is added to the routing table.
    - If you are changing an existing route, click the **Apply** button. Your settings are saved.
11. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 116.Learned routes information

Field	Description
Network Address	The IP route prefix for the destination.
Subnet Mask	The subnet/network mask, which indicates the portion of the IP interface address that identifies the attached network.
Protocol	The protocol that created the route: Local, Static, OSPF, or RIP.
Route Type	The type of route, depending on the protocol: Connected, Static, or Dynamic.
Next Hop Interface	The interface of the outgoing router interface that must be used when traffic is forwarded to the destination.
Next Hop Address	The IP address of the outgoing router interface that must be used when traffic is forwarded to the next router (if any) in the path toward the destination.
Preference	The preference, which is a value from 0 to 255.
Metric	The administrative cost of the path to the destination. The default is 1. The range is from 0 to 255.

Delete a route

You can delete a route that you no longer need.

To delete a route:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing> Routing Table > Basic > Route Configuration**.
   The Route Configuration page displays.
3. Select the check box for the route.
4. Click the **Delete** button.
   Your settings are saved. The route is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Specify route preferences

You can configure the default preference for each protocol, for example, 60 for static routes or 120 for RIP. These values are arbitrary values in the range from 1 to 255 and are independent of route metrics. Most routing protocols use a route metric to determine the shortest path known to the protocol, independent of any other protocol.
The best route to a destination is chosen by selecting the route with the lowest preference value. If multiple routes exist to a destination, the preference values are used to determine the preferred route. If there is still a tie, the route with the best route metric is chosen. To prevent problems with mismatched metrics such as protocols that are not directly comparable, you must configure different preference values for each of the protocols.

To specify route preferences

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Routing Table > Advanced > Route Preferences**.
   The Route Preferences page displays.
3. In the **Static** field, specify the static route preference on the switch.
   The default value is 1. The range is from 1 to 255.
4. In the **RIP** field, specify the RIP route preference on the switch.
   The default value is 120. The range is from 1 to 255.
5. In the **OSPF Intra** field, specify the OSPF intra route preference on the switch.
   The default value is 110. The range is from 1 to 255.
6. In the **OSPF Inter** field, specify the OSPF inter route preference on the switch.
   The default value is 110. The range is from 1 to 255.
7. In the **OSPF External** field, specify the OSPF external route preference on the switch.
   The default value is 110. The range is from 1 to 255.
8. Click the **Apply** button.
   Your settings are saved.
   The Local field displays the local route preference value.
9. To save the settings to the running configuration, click the **Save** icon.

IPv4 routing

You can enable or disable the IPv4 routing mode, configure the global IPv4 routing settings, configure IPv4 routing interfaces, add a secondary IP address to an IPv4 routing interface, and

view IPv4 routing statistics.

Manage the global IPv4 routing settings

You can configure the IPv4 routing settings for the switch, as opposed to the IPv4 routing settings for an interface.

To configure the global IPv4 routing settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IP > Basic > IP Configuration**.
   The IP Configuration page displays.
3. Select the Routing Mode **Enable** or **Disable** radio button.
   This selection enables or disables the administrative mode for IPv4 routing for the switch. The default is Enable.
4. Select the ICMP Echo Replies **Enable** or **Disable** radio button.
   The default is Enable, which means that the switch sends ICMP echo replies to ICMP echo requests.
5. Select the ICMP Redirects **Enable** or **Disable** radio button.
   The default is Disable. If you enable ICMP Redirects, the switch can forward ICMP Redirects.
6. In the **ICMP Rate Limit Interval** field, specify the number of ICMP error packets that are allowed per burst interval.
   By default, the burst interval is 1000 msec (the rate limit is 100 packets/sec). The interval is from 0 to 2147483647 msec. Enter 0 to let the switch ignore the number of ICMP error packets that are allowed per burst interval.
7. In the **ICMP Rate Limit Burst Size** field, specify the number of ICMP error packets that are allowed for the burst size.
   By default, the burst size is 100 packets. The burst size range is from 1 to 200. Enter 0 to let the switch ignore the number of ICMP error packets that are allowed for the burst size.
8. To configure a global default gateway, select the **Select to configure Global Default Gateway** check box. Then, in the **Global Default Gateway** field, enter the IP address.
   The default gateway that you specify receives a higher preference than a default gateway that is learned from a DHCP server. You configure one default gateway only, but you can change the default gateway.
9. Click the **Apply** button.
   Your settings are saved.

10. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 117.IPv4 routing configuration information

Field	Description
Default Time to Live	The default value inserted into the Time-To-Live field of the IP header of datagrams originated by the switch, if a TTL value is not supplied by the transport layer protocol.
Maximum Next Hops	The maximum number of hops supported by the switch.
Maximum Routes	The maximum number of routes supported by the switch.
Maximum Static Routes	The maximum number of static routes supported by the switch.

Display the IPv4 statistics

You can display the IPv4 routing statics for the switch.

To display the IPv4 routing statistics for the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.

The System Information page displays.

2. Select **Routing > IP > Basic > Statistics**.
   The Statistics page displays.
   The following table describes the view-only fields on the page.

Table 118.IP statistics information

Field	Description
IpInReceives	The total number of input datagrams received from interfaces, including those received in error.
IpInHdrErrors	The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, and so on.
IpInAddrErrors	The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (for example, 0.0.0.0) and addresses of unsupported classes (Class E). For entities that are not IP gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
IpForwDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities that do not act as IP gateways, this counter includes only those packets that were source-routed through this entity, and the source-route option processing was successful.
IpInUnknownProtos	The number of locally addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
IpInDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but that were discarded (for lack of buffer space). This counter does not include any datagrams discarded while awaiting reassembly.
IpInDelivers	The total number of input datagrams successfully delivered to IP user protocols (including ICMP).
IpOutRequests	The total number of IP datagrams that local IP user protocols (including ICMP) supplied to IP in requests for transmission. This counter does not include any datagrams counted in ipForwDatagrams.
IpOutDiscards	The number of output IP datagrams for which no problem was

Field	Description
	encountered to prevent their transmission to their destination, but that were discarded for reasons such as lack of buffer space. This counter would include datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
IpOutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. This counter includes any packets counted in ipForwDatagrams that meet this no-route criterion. This includes any datagrams that a host cannot route because all of its default gateways are down.
IpReasmTimeout	The maximum number of seconds for which received fragments are held while they are awaiting reassembly at this entity.
IpReasmReqds	The number of IP fragments received that were reassembled at this entity.
IpReasmOKs	The number of IP datagrams successfully reassembled.
IpReasmFails	The number of failures detected by the IP reassembly algorithm (for whatever reason: timed out, errors, and so on). This is not necessarily a count of discarded IP fragments since some algorithms can lose track of the number of fragments by combining them as they are received.
IpFragOKs	The number of IP datagrams that were fragmented at this entity.
IpFragFails	The number of IP datagrams that were discarded because they needed to be fragmented at this entity but could not be, for reasons such as their Don't Fragment flag was set.
IpFragCreates	The number of IP datagram fragments that were generated as a result of fragmentation at this entity.
IpRoutingDiscards	The number of routing entries that were discarded even though they were valid. One possible reason for discarding such an entry could be to free up buffer space for other routing entries.
IcmpInMsgs	The total number of ICMP messages that the entity received. This counter includes all those counted by icmpInErrors.
IcmpInErrors	The number of ICMP messages that the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, and so on).
IcmpInDestUnreachs	The number of ICMP destination unreachable messages received.
IcmpInTimeExcds	The number of ICMP time exceeded messages received.

Field	Description
IcmpInParmProbs	The number of ICMP parameter problem messages received.
IcmpInSrcQuenchs	The number of ICMP source quench messages received.
IcmpInRedirects	The number of ICMP redirect messages received.
IcmpInEchos	The number of ICMP echo (request) messages received.
IcmpInEchoReps	The number of ICMP echo reply messages received.
IcmpInTimestamps	The number of ICMP timestamp (request) messages received.
IcmpInTimestampReps	The number of ICMP timestamp reply messages received.
IcmpInAddrMasks	The number of ICMP address mask request messages received.
IcmpInAddrMaskReps	The number of ICMP address mask reply messages received.
IcmpOutMsgs	The total number of ICMP messages that this entity attempted to send. This counter includes all those counted by icmpOutErrors.
IcmpOutErrors	The number of ICMP messages that this entity did not send due to problems discovered within ICMP such as a lack of buffers. This value does not include errors discovered outside the ICMP layer such as the inability of IP to route the resultant datagram. In some implementations there might be no types of error that contribute to this counter's value.
IcmpOutDestUnreachs	The number of ICMP destination unreachable messages sent.
IcmpOutTimeExcds	The number of ICMP time exceeded messages sent.
IcmpOutParmProbs	The number of ICMP parameter problem messages sent.
IcmpOutSrcQuenchs	The number of ICMP source quench messages sent.
IcmpOutRedirects	The number of ICMP redirect messages sent. For a host, this is always zero, since hosts do not send redirects.
IcmpOutEchos	The number of ICMP echo (request) messages sent.
IcmpOutEchoReps	The number of ICMP echo reply messages sent.
IcmpOutTimestamps	The number of ICMP timestamp (request) messages.
IcmpOutTimestampReps	The number of ICMP timestamp reply messages sent.
IcmpOutAddrMasks	The number of ICMP address mask request messages sent.

Configure IPv4 routing interfaces

You can configure one or more IPv4 routing interfaces.

To configure one or more IPv4 routing interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing> IP > Advanced > IP Interface Configuration**.
   The IP Interface Configuration page displays.
3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. In the **Description** field, enter a description for the interface.
6. From the **IP Address Configuration Method** menu, select the method by which an IP address is configured on the interface: **None**, **Manual**, or **DHCP**.
   By default, the method is None. Use the None method to reset the DHCP method.

NOTE: When you change the configuration method from DHCP to None, a minor delay occurs before the page refreshes.

7. If you select **Manual** from the **IP Address Configuration Method** menu, do the following:
   a. In the **IP Address** field, enter the IP address for the interface.
   b. In the **Subnet Mask**, enter the subnet mask for the interface.
      This is also referred to as the subnet/network mask, and defines the portion of the interface's IP address that is used to identify the attached network.
8. From the **Routing Mode** menu, select **Enable** or **Disable** to enable or disable the routing mode for the interface.
   The default is Enable.
9. From the **Administrative Mode** menu, select **Enable** or **Disable** to enable or disable the administrative mode of the interface.
   The default is Enable.
10. From the **Forward Net Directed Broadcasts** menu, select **Enable** or **Disable** to enable or disable the forwarding of network-directed broadcast packets.
    The default is Disable, which means that network-directed broadcast packets are not forwarded but dropped.
11. From the**Encapsulation Type** menu, select the link layer encapsulation type for packets transmitted from the specified interface.
    You can select **Ethernet** or **SNAP**, which is Subnetwork Access Protocol. The default is Ethernet.
12. From the **Proxy Arp** menu, select **Enable** or **Disable** to enable or disable the proxy ARP for the interface.
    The default is Enable.
13. From the **Local Proxy Arp** menu, select **Enable** or **Disable** to enable or disable the local proxy ARP for the interface.
    The default is Disable.
14. In the **Bandwidth (kbps)** field, specify the bandwidth for the interface.
    This information in this field determines the speed of the interface to higher level protocols. Some protocols use bandwidth to compute the link cost. The range is from 1 to 10000000.
15. From the **ICMP Destination Unreachables** menu, select **Enable** or **Disable** to enable or disable the interface from sending ICMP destination unreachables messages.
    The default is Enable.
16. From the **ICMP Redirects** menu, select **Enable** or **Disable** to enable or disable the ICMP redirecting mode.
    The interface sends an ICMP redirect message only if this function is enabled both globally and on the interface. The default is Enable.

17. In the **IP MTU** field, specify the maximum size of IP packets that are sent on the interface. The IP MTU is the maximum frame size minus the length of the Layer 2 header. The range is from 68 bytes to the MTU size that is supported by the interface link. The default is 1500. A value of 0 indicates that the IP MTU is not configured, in which case the interface uses the MTU size that is supported by the interface link.
18. Click the **Apply** button.
    Your settings are saved.
19. To save the settings to the running configuration, click the **Save** icon.
    The following table describes the view-only fields on the page.

Table 119.IP Interface Configuration

Field	Description
VLAN ID	The VLAN ID for the interface
Link Speed Data Rate	The physical link data rate of the interface
OSPF Admin Mode	Indicates if OSPF is enabled (Enable) or disabled (Disable) for the interface
Active State	The active or inactive state of the interface (Active or Inactive)
MAC Address	The MAC address that is associated with the interface
Link State	The link-up or link-down state of the interface. (Link Up or Link Down.) An interface is considered active if the link is up and the interface is in the forwarding state. An interface that is configured for stacking in the detached state (Detach).
Routing Interface Status	The link status of the routing interface (Up or Down)

Delete the routing IP address from an IPv4 routing interface

You can delete the routing IP address from one or more IPv4 routing interfaces.

To delete the routing IP address from one or more IPv4 routing interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IP > Advanced > IP Interface Configuration**.
   The IP Interface Configuration page displays.

3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To delete the IP address for a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To delete the IP address for multiple interfaces, select the check box associated with each interface.
5. Click the **Delete** button.
   Your settings are saved. The IP routing address is removed from the IPv4 routing interface.
6. To save the settings to the running configuration, click the **Save** icon.

Configure a secondary IP address for an IPv4 routing interface

You can configure a secondary IP address for an IPv4 routing interface.

To configure the secondary IP address:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IP > Advanced > Secondary IP**.
   The Secondary IP page displays.
3. From the **Interface** menu, select the interface.
4. In the **Secondary IP Address** field, add a secondary IPv4 address to the interface.
5. In the **Secondary IP Subnet Mask** field, enter the subnet mask associated with the secondary IP address.
6. Click the **Add** button.
   The secondary IP address is added.

7. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 120.Secondary IP address information

Field	Description
VLAN ID	The VLAN ID associated with the selected interface.
Primary IP Address	The primary IP address for the interface.

Delete the secondary IP address from an IPv4 routing interface

You can delete the secondary IP address from an IPv4 routing interface.

To delete the secondary IP address from an IPv4 routing interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IP > Advanced > Secondary IP**.
   The Secondary IP page displays.
3. From the **Interface** menu, select the interface.
4. In the Secondary IP Address table, select the check box for the IP address.
5. Click the **Delete** button.
   Your settings are saved. The secondary IP routing address is removed from the IPv4 routing interface.
6. To save the settings to the running configuration, click the **Save** icon.

IPv6 routing

You can enable or disable the IPv6 routing mode, configure the global IPv6 routing settings, configure IPv6 routing interfaces, add IPv6 prefixes, add IPv6 static routes and route preferences, add IPv6 tunnels, and view IPv6 routing statistics, routes, and neighbors.

Manage the global IPv6 routing settings

You can configure the IPv6 routing settings for the switch, as opposed to the IPv6 routing settings for an interface.

To configure the global IPv6 routing settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 > Basic > Global Configuration**.
   The Global Configuration page displays.
3. Select the IPv6 Unicast Routing **Enable** or **Disable** radio button.
   This selection enables or disables the administrative mode for IPv6 routing for the switch. The default is Disable.
4. In the **Hop Limit** field, enter a value for the unicast hop count that is used for IPv6 packets that originate from the switch.
   The hop count is also included in router advertisements, and ranges from 1 to 255, inclusive. The default is 64.
5. In the **ICMPv6 Rate Limit Error Interval** field, specify the number of ICMP error packets that are allowed per burst interval.
   By default, the burst interval is 1000 msec (the rate limit is 100 packets/sec). The interval is from 0 to 2147483647 msec. Enter 0 to let the switch ignore the number of ICMPv6 error packets that are allowed per burst interval.
6. In the **ICMPv6 Rate Limit Burst Size** field, specify the number of ICMP error packets that are allowed for the burst size.
   By default, the burst size is 100 packets. The burst size range is from 1 to 200. Enter 0 to let the switch ignore the number of ICMPv6 error packets that are allowed for the burst size.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Display the IPv6 route table

The IPv6 route table includes all routes, including the best routes and the configured routes. You can select which routes to display.

To display the IPv6 route table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 > Basic > Route Table**.
   The Route Table page displays.

3. From the **Routes Displayed** menu, select which routes must be displayed in the table:
   - **All Routes**: All active IPv6 routes.
   - **Best Routes Only**: The best active routes only,
   - **Configured Routes Only**: The manually configured routes only.
4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 121.IPv6 route table information

Field	Description
Number of Routes	The total number of active routes in the route table.
IPv6 Prefix	The network prefix for the active route.
Prefix Length	The prefix length for the active route.
Protocol	The type of protocol for the active route.
Next Hop Interface	The interface over which the route is active. For a reject route, the next hop is a null (Null0) interface.
Next Hop IP Address	The next hop IPv6 address for the active route.
Preference	The route preference of the configured route.

Configure IPv6 routing interfaces

You can configure one or more IPv6 routing interfaces.

To configure one or more IPv6 routing interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 > Advanced > Interface Configuration**.
   The Interface Configuration page displays.

3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **IPv6 Mode** menu, select **Enable** or **Disable** to enable or disable the IPv6 routing mode for the interface.
   When the IPv6 routing mode is enabled, the interface is capable of IPv6 operation without a global address. In this case, an EUI-64 based link-local address is used. The default is Disable.
6. From the **DHCPv6 Client Mode** menu, select **Enable** or **Disable** to enable or disable DHCPv6 client mode on the interface.
   Only one interface can function as a DHCPv6 client. The default is Disable.
7. From the **Stateless Address AutoConfig Mode** menu, select **Enable** or **Disable** to enable or disable the stateless address automatic configuration mode on the interface.
   The default value is Disable.
8. From the **Routing Mode** menu, select **Enable** or **Disable** to enable or disable the routing mode for the interface.
   The default is Disable.
9. From the **Admin Mode** list, select **Enable** or **Disable** to enable or disable the administrative mode of the interface.
   The default is Enable.
10. In the **MTU** field, specify the maximum transmit unit for the interface.
    The range is from 1280 to 1500. The default is 1500. If you enter is 0, you disable routing on the interface.

11. In the **Duplicate Address Detection Transmits** field, specify the number of duplicate address detection (DAD) transmits for the interface.
    The DAD transmits value must be in the range from 0 to 600. The default is 1.
12. In the **Life Time Interval** field, specify the lifetime interval that the interface transmits in router advertisements.
    This is the information about the period during which IPv6 neighbors can use the interface as a default router.
    The range is from 0 to 9000. The default is 1800. The value must be greater than or equal to the maximum advertisement interval. If you enter 0, the interface cannot be used as the default routing interface.
13. In the **Adv NS Interval** field, specify the retransmission interval information that the interface transmits in router advertisements.
    This is the information about the interval between neighbor solicitation retransmissions.
    The range is from 1000 to 4294967295. The default is 0.
14. In the **Adv Reachable Interval** field, specify the reachable interval information that the interface transmits in router advertisements.
    This is the information about the period during which a remote IPv6 device is considered reachable after receipt of a neighbor discovery confirmation message.
    The range is from 0 to 3600000. The default is 0.
15. In the **Adv Interval** field, specify the period between router advertisements from the interface.
    The range is from 4 to 1800. The default is 600.
16. From the **Adv Management Config Flag** menu, select **Enable** or **Disable** to enable or disable the router advertisement “managed address configuration flag” for the interface.
    When enabled, end nodes use DHCPv6. When disabled, end nodes automatically configure addresses. The default is Disable.
17. From the **Adv Other Config Flag** menu, select **Enable** or **Disable** to enable or disable the router advertisement “other stateful configuration flag” for the interface.
    The default is Disable. The other stateful configuration flag informs an IPv6 host to use DHCPv6 to get additional configuration information such as, for example, the IPv6 address of a DNS server.
18. From the **Router Preference** menu, select the router preference advertisements for the interface.
    You can select **High**, **Medium**, or **Low**. The default is Medium.

19. In the **Adv Suppress Flag** list, select **Enable** or **Disable** to enable or disable the router advertisement suppression for the interface.
    The default is Disable.
20. From the **Destination Unreachables** menu, select **Enable** or **Disable** to enable or disable the transmission of ICMPv6 destination unreachables message on the interface.
    The default is Enable.
21. Click the **Apply** button.
    Your settings are saved.
22. To save the settings to the running configuration, click the **Save** icon.
    The following table describes the view-only fields on the page.

Table 122.IPv6 interface configuration information

Field	Description
Operational Mode	The operational state of an interface. The default is Disable.
Link State	Indicates if the link is up or down (Link Up or Link Down).

Configure prefix settings for an IPv6 routing interface

You can add or change one or more prefixes and associated settings for an IPv6 routing interface.

To add or change prefix settings for an IPv6 routing interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 > Advanced > Prefix Configuration**.
   The Prefix Configuration page displays.
3. From the **Interface** menu, select the interface.
4. If you are changing en existing prefix for the interface, select the check box for the prefix.
5. If you are adding a new prefix, in the **IPv6 Prefix** field, specify the IPv6 prefix for the interface.
6. If you are adding a new prefix, in the **Prefix Length** field, specify the IPv6 prefix length for the interface.
7. From the **EUI64** menu, select **Enable** or **Disable** to enable or disable the 64-bit unicast prefix.

8. In the **Valid Life Time** field, specify the valid period for prefix router advertisements for the interface.
   This is the period during which the prefix is valid for the purpose of on-link determination. The range is from 0 to 4294967295.
9. In the **Preferred Life Time** field, specify preferred period for prefix router advertisements for the interface.
   This is the period during which the prefix is preferred for the purpose of on-link determination. An autoconfigured address generated from the prefix is a preferred address. The range from 0 to 4294967295.
10. From the **Onlink Flag** menu, select **Enable** or **Disable** to enable or disable the prefix from being used for autonomous address configuration.
    If enabled, the prefix can be used for on-link determination.
11. From the **Autonomous Flag** menu, to enable or disable the prefix from being used for on-link determination.
    If enabled, the prefix can be used for autonomous address configuration.
12. Do one of the following:
    - If you are adding a new prefix, click the **Add** button. Your settings are saved. The prefix is added for the interface.
    - If you are changing an existing prefix, click the **Apply** button. Your settings are saved.

    The Current State field displays the state of the IPV6 address. The state is TENT (tentative) if routing is disabled or DAD (Duplicate Address Detection) fails. The state is Active if the interface is active and DAD is successful.
13. To save the settings to the running configuration, click the **Save** icon.

Delete a prefix setting from an IPv6 routing interface

You can delete a prefix and associated settings that you no longer need for an IPv6 routing interface.

To a delete prefix settings from an IPv6 routing interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 > Advanced > Prefix Configuration**.
   The Prefix Configuration page displays.
3. From the **Interface** menu, select the interface.
4. Select the check box for the prefix.

5. Click the **Delete** button.
   Your settings are saved. The prefix is removed.
6. To save the settings to the running configuration, click the **Save** icon.

Display the IPv6 and ICMPv6 statistics for an IPv6 routing interface

You can display the IPv6 and ICMPv6 statics for an IPv6 routing interface.

To display the IPv6 and ICMPv6 statistics for an IPv6 routing interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 > Advanced > Statistics**.
   The Statistics page displays.
3. From the **Interface** menu, select the IPv6 routing interface.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields in the IPv6 Statistics section.

Table 123.IPv6 interface statistics information

Field	Description
Total Datagrams Received	The total number of input datagrams received by the interface, including those received in error
Received Datagrams Locally Delivered	The total number of datagrams successfully delivered to IPv6 user-protocols (including ICMP). This counter is incremented at the interface to which these datagrams were addressed, which might not be the input interface for some of the datagrams.
Received Datagrams Discarded Due To Header Errors	The number of input datagrams discarded due to errors in their IPv6 headers, including version number mismatch, other format errors, hop count exceeded, errors discovered in processing their IPv6 options, and so on
Received Datagrams Discarded Due To MTU	The number of input datagrams that could not be forwarded because their size exceeded the link MTU of outgoing interface
Received Datagrams Discarded Due To No Route	The number of input datagrams discarded because no route could be found to transmit them to their destination
Received Datagrams With Unknown Protocol	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol. This counter is incremented at the interface to which these datagrams were addressed, which might not be the input interface for some of the datagrams.
Received Datagrams Discarded Due To Invalid Address	The number of input datagrams discarded because the IPv6 address in their IPv6 header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (for example, ::0) and unsupported addresses (such as addresses with unallocated prefixes). For entities that are not IPv6 routers and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
Received Datagrams Discarded Due To Truncated Data	The number of input datagrams discarded because datagram frame did not carry enough data
Received Datagrams Discarded Other	The number of input IPv6 datagrams for which no problems were encountered to prevent their continued processing, but that were discarded for reasons such as lack of buffer space. This counter does not include any datagrams discarded while awaiting re-

Field	Description
	assembly.
Received Datagrams Reassembly Required	The number of IPv6 fragments received that needed to be reassembled at this interface. This counter is incremented at the interface to which these fragments were addressed, which might not be the input interface for some of the fragments.
Datagrams Reassembled	The number of IPv6 datagrams successfully reassembled. This counter is incremented at the interface to which these datagrams were addressed, which might not be necessarily the input interface for some of the fragments.
Datagrams Failed To Reassemble	The number of failures detected by the IPv6 reassembly algorithm (for whatever reason: timed out, errors, and so on). This is not necessarily a count of discarded IPv6 fragments since some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received. This counter is incremented at the interface to which these fragments were addressed, which might not be the input interface for some of the fragments.
Datagrams Forwarded	The number of output datagrams that this entity received and forwarded to their final destinations. In entities that do not act as IPv6 routers, this counter includes only those packets that were source-routed through this entity, and the source-route processing was successful. For a successfully forwarded datagram the counter of the outgoing interface is incremented.
Datagrams Locally Transmitted	The number of datagrams that this entity successfully transmitted from this output interface.
Datagrams Transmit Failed	The number of datagrams that this entity failed to transmit successfully
Datagrams Fragmented	The number of IPv6 datagrams that were fragmented at this output interface
Datagrams Failed To Fragment	The number of output datagrams that could not be fragmented at this interface
Datagrams Fragments Created	The number of output datagram fragments that were generated as a result of fragmentation at this output interface
Multicast Datagrams Received	The number of multicast packets received by the interface
Multicast Datagrams Transmitted	The number of multicast packets transmitted by the interface

The following table describes the view-only fields in the ICMPv6 Statistics section.

Table 124.ICMPv6 interface statistics information

Field	Description
Total ICMPv6 Messages Received	The total number of ICMP messages received by the interface, which includes all those counted by IPv6IfIcmpInErrors. This interface is the interface to which the ICMP messages were addressed, which might not be the input interface for the messages.
ICMPv6 Messages With Errors Received	The number of ICMP messages that the interface received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, and so on)
ICMPv6 Destination Unreachable Messages Received	The number of ICMP Destination Unreachable messages received by the interface
ICMPv6 Messages Prohibited Administratively Received	The number of ICMP destination unreachable/communication administratively prohibited messages received by the interface
ICMPv6 Time Exceeded Messages Received	The number of ICMP Time Exceeded messages received by the interface
ICMPv6 Parameter Problem Messages Received	The number of ICMP Parameter Problem messages received by the interface
ICMPv6 Packet Too Big Messages Received	The number of ICMP Packet Too Big messages received by the interface
ICMPv6 Echo Request Messages Received	The number of ICMP Echo (request) messages received by the interface
ICMPv6 Echo Reply Messages Received	The number of ICMP Echo Reply messages received by the interface
ICMPv6 Router Solicit Messages Received	The number of ICMP Router Solicit messages received by the interface
ICMPv6 Router Advertisement Messages Received	The number of ICMP Router Advertisement messages received by the interface
ICMPv6 Neighbor Solicit Messages Received	The number of ICMP Neighbor Solicit messages received by the interface
ICMPv6 Neighbor Advertisement Messages Received	The number of ICMP Neighbor Advertisement messages received by the interface

Field	Description
ICMPv6 Redirect Messages Received	The number of ICMPv6 Redirect messaged received by the interface
ICMPv6 Group Membership Query Messages Received	The number of ICMPv6 Group Membership Query messages received by the interface
ICMPv6 Group Membership Response Messages Received	The number of ICMPv6 Group Membership Response messages received by the interface
ICMPv6 Group Membership Reduction Messages Received	The number of ICMPv6 Group Membership Reduction messages received by the interface
Total ICMPv6 Messages Transmitted	The total number of ICMP messages that this interface attempted to send. This counter includes all those counted by icmpOutErrors.
ICMPv6 Messages Not Transmitted Due To Error	The number of ICMP messages that this interface did not send due to problems discovered within ICMP such as a lack of buffers. This value does not include errors discovered outside the ICMP layer such as the inability of IPv6 to route the resultant datagram. In some implementations there might be no types of error that contribute to this counter's value.
ICMPv6 Destination Unreachable Messages Transmitted	The number of ICMP Destination Unreachable messages sent by the interface
ICMPv6 Messages Prohibited Administratively Transmitted	Number of ICMP Destination Unreachable/Communication Administratively Prohibited messages sent
ICMPv6 Time Exceeded Messages Transmitted	The number of ICMP Time Exceeded messages sent by the interface
ICMPv6 Parameter Problem Messages Transmitted	The number of ICMP Parameter Problem messages sent by the interface
ICMPv6 Packet Too Big Messages Transmitted	The number of ICMP Packet Too Big messages sent by the interface
ICMPv6 Echo Request Messages Transmitted	The number of ICMP Echo (request) messages sent by the interface
ICMPv6 Echo Reply Messages Transmitted	The number of ICMP Echo Reply messages sent by the interface
ICMPv6 Router Solicit Messages Transmitted	The number of ICMP Neighbor Solicitation messages sent by the interface
ICMPv6 Router Advertisement	The number of ICMP Router Advertisement messages sent by the

Field	Description
Messages Transmitted	interface
ICMPv6 Neighbor Solicit Messages Transmitted	The number of ICMP Neighbor Solicitation messages sent by the interface
ICMPv6 Neighbor Advertisement Messages Transmitted	The number of ICMP Neighbor Advertisement messages sent by the interface
ICMPv6 Redirect Messages Transmitted	The number of Redirect messages sent
ICMPv6 Group Membership Query Messages Transmitted	The number of ICMPv6 Group Membership Query messages sent
ICMPv6 Group Membership Response Messages Transmitted	The number of ICMPv6 Group Membership Response messages sent
ICMPv6 Group Membership Reduction Messages Transmitted	The number of ICMPv6 Group Membership Reduction messages sent
ICMPv6 Duplicate Address Detects	The number of duplicate addresses detected by the interface

Display the IPv6 neighbor table or clear IPv6 neighbor entries

You can display the IPv6 neighbor devices that the switch detects or clear IPv6 neighbor entries from the neighbor table.

To display the IPv6 neighbor table or clear IPv6 neighbor entries:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 > Advanced > Neighbor Table**.
   The Neighbor Table page displays.
3. From the **Search By** menu, select one of the following methods to search for IPv6 neighbors:
   - **IPv6 Address**: Enter the 128-byte hexadecimal IPv6 address in four-digit groups separated by colons, for example, 2001:231F:::1. Then click the **Go** button. If the address exists, that entry is displayed. An exact match is required.
   - **Interface**: Enter the interface ID in unit/port format, for example, 0/4. Then click the **Go** button. If the address exists, that entry is displayed.

4. To clear the IPv6 neighbors that display, click the **Clear** button.
5. To save the settings to the running configuration, click the **Save** icon.
6. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 125.IPv6 Advanced Neighbor Table

Field	Description
Interface	The interface for which settings are displayed
IPv6 Address	The IPv6 address of the neighbor or interface
MAC Address	The MAC address associated with an interface
isRtr	Displays if the neighbor is a router. If the neighbor is a router, the field displays True. If the neighbor is not a router, the field displays False.
Neighbor State	The state of the neighbor cache entry. Following are the states for dynamic entries in the IPv6 neighbor discovery cache: ○ **Incmp**: Address resolution is being performed on the entry. A neighbor solicitation message was sent to the solicited-node multicast address of the target, but the corresponding neighbor advertisement message was not yet received. ○ **Reach**: Positive confirmation was received within the last "ReachableTime" milliseconds that the forward path to the neighbor was functioning properly. While in REACH state, the device takes no special action as packets are sent. ○ **Stale**: More than the "ReachableTime" milliseconds elapsed since the last positive confirmation was received that the forward path was functioning properly. While in STALE state, the device takes no action until a packet is sent. ○ **Delay**: More than "ReachableTime" milliseconds elapsed since the last positive confirmation was received that the forward path was functioning properly. A packet was sent within the last DELAY_FIRST_PROBE_TIME seconds. If no reachability confirmation is received within DELAY_FIRST_PROBE_TIME seconds of entering the DELAY state, send a neighbor solicitation message and change the state to PROBE. ○ **Probe**: A reachability confirmation is actively sought by resending neighbor solicitation messages every "RetransTimer" milliseconds until a reachability confirmation is received.
Last Updated	The time since the address was confirmed to be reachable

Configure IPv6 static routes

You can add or change a static IPv6 route.

To add or change a static IPv6 route:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 > Advanced > Static Route Configuration**.
   The Static Route Configuration page displays.
3. If you are changing en existing route, select the check box for the route.
   The only field you can change is the preference for the route.
4. If you are adding a new route, in the **IPv6 Prefix** field, specify the IPv6 prefix for the route.
5. If you are adding a new route, in the **Prefix Length** field, specify the IPv6 prefix length for the route.
6. If you are adding a new route, from the **Next Hop IPv6 Address Type** menu, select one of the following options:
   - **Global**: Create a global IPv6 route by specifying the interface for the route.
   - **Link-Local**: Create a link-local IPv6 address by specifying the interface and the next hop address for the route.
   - **Static-Reject**: Create a static-reject route for a destination prefix. You cannot specify an interface or a next hop address for the route.
7. If you are adding a new route and selected **Global** or **Link-Local** from the **Next Hop IPv6 Address Type** menu, specify the IPv6 address for the next hop in the **Next Hop IPv6 Address** field.
8. If you are adding a new route and selected **Link-Local** from the **Next Hop IPv6 Address Type** menu, specify the interface for the route.
9. In the **Preference** field, specify the preference for the route.
   The range is from 1 to 255. You can add the preference for a new route or change the preference for an existing route.
10. Do one of the following:
    - If you are adding a new route, click the **Add** button. Your settings are saved. The route is added.
    - If you are changing an existing route, click the **Apply** button. Your settings are saved.
11. To save the settings to the running configuration, click the **Save** icon.

Delete an IPv6 static route

You can delete an IPv6 static route that you no longer need.

To delete an IPv6 static route:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 > Advanced > Static Route Configuration**.
   The Static Route Configuration page displays.
3. Select the check box for the route.
4. Click the **Delete** button.
   Your settings are saved. The route is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the IPv6 route preference for the switch

You can configure the default preference for the switch, which is a value in the range from 1 to 255 and is independent of route metrics. Most routing protocols use a route metric to determine the shortest path known to the protocol, independent of any other protocol. The best route to a destination is automatically selected by determining the route with the lowest preference value. If multiple routes to a destination exist, the preference value is used to determine the preferred route. If a tie occurs, the route with the best route metric is selected.

Configure the IPv6 route preference for the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 > Advanced > Route Preference**.
   The Route Preference page displays.
3. In the **Static** field, specify the static route preference value for the switch
   The range is from 1 to 255. The default value is 1.
4. In the **OSPFv3 Intra** field, specify the OSPFv3 intra route preference on the switch.
   The default value is 110. The range is from 1 to 255.
5. In the **OSPFv3 Inter** field, specify the OSPFv3 inter route preference on the switch.
   The default value is 110. The range is from 1 to 255.

6. In the **OSPFv3 External** field, specify the OSPFv3 external route preference on the switch. The default value is 110. The range is from 1 to 255.
7. Click the **Apply** button.
   Your settings are saved.
   The Local field displays the local preference.
8. To save the settings to the running configuration, click the **Save** icon.

Routing VLANs

You can configure some interfaces to support VLANs and other interfaces to support routing. You can also configure a VLAN to function as if it were a routing interface.
When an interface is enabled for bridging (the default setting) rather than routing, all normal bridge processing is performed for an inbound packet, which is then associated with a VLAN. Its MAC destination address (MAC DA) and VLAN ID are used to search the MAC address table. If routing is enabled for the VLAN, and the MAC DA of an inbound unicast packet is that of the internal bridge-router interface, the packet is routed. An inbound multicast packet is forwarded to all ports in the VLAN, plus the internal bridge-router interface, if it was received on a routed VLAN.
Because an interface can belong to more than one VLAN, VLAN routing might be enabled for all of the VLANs on the interface, or for a subset. VLAN routing can be used to allow more than one physical port to reside on the same subnet. It could also be used when a VLAN spans multiple physical networks, or when additional segmentation or security is required.
An interface can be either a VLAN interface or a routing interface, but not both. However, a VLAN interface can be part of a VLAN that functions as a routing interface.

Create a routing VLAN with the VLAN static routing wizard

The VLAN static routing wizard lets you create a VLAN and add interfaces to the VLAN. The VLAN static routing wizard also lets you add selected interfaces as a link aggregation group (LAG). With the wizard, you can do the following:

- Create a VLAN and generate a unique name for VLAN.
- Add selected interfaces to the newly created VLAN or remove selected interfaces from the default VLAN.
- Add selected interfaces to a LAG and add the LAG to a newly created VLAN.
- Enable tagging on selected interfaces if the interface is a member of another VLAN. Or, disable tagging if the selected interface is not a member of another VLAN.
- Exclude ports from the VLAN.
- Enable routing on the VLAN and specify an IP address and subnet mask for routing.

To create a routing VLAN with the VLAN static routing wizard:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > VLAN > VLAN Static Routing Wizard**.
   The VLAN Static Routing Wizard page displays.
   The ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed. In addition, all LAGs are displayed, whether or not a stack is configured.
3. In the **VLAN ID** field, specify the VLAN identifier (VID) that must be associated with the VLAN.
   The range is from 1 to 4093.
4. In the **IP Address** field, specify the IPv4 address for the VLAN routing interface.
5. In the **Network Mask** field, specify the subnet mask for the VLAN routing interface.
6. In the Ports table (or if a stack is configured, in one of the Ports tables), click a port once, twice, or three times to configure one of the following modes:
   - **T (tagged) member**: The port is added as a tagged member of the VLAN.
   - **U (untagged) member**: The port is added as an untagged member of the VLAN.
   - **Excluded member**: By default, the selection is blank, which means that the port is excluded from the VLAN. A port that is excluded can still be dynamically added to the VLAN through GVRP.

   If a stack is configured, you can select ports in multiple tables.
7. In the LAG table, click a LAG once, twice, or three times to configure one of the following modes:
   - **T (tagged) member**: The LAG is added as a tagged member of the VLAN.
   - **U (untagged) member**: The LAG is added as an untagged member of the VLAN.
   - **Excluded member**: By default, the selection is blank, which means that the LAG is excluded from the VLAN. A LAG that is excluded can still be dynamically added to the VLAN through GVRP.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Configure routing for an existing VLAN

For an existing VLAN, you can add or change an IPv4 address and subnet mask and configure routing for the VLAN.

To add or change routing for an existing VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > VLAN > VLAN Routing Configuration**.
   The VLAN Routing page displays.
3. If you are changing an existing routing VLAN, select the check box for the VLAN.
4. If you are adding a new routing VLAN, from the **VLAN ID** menu, select the VLAN ID.
   This menu displays the IDs of all VLANs that are configured on this switch.
5. In the **IP Address** field, specify the IP address for the routing VLAN.
6. In the **Subnet Mask** field, specify the subnet mask for the routing VLAN.
7. Do one of the following:
   - If you are adding a new routing VLAN, click the **Add** button. Your settings are saved. The routing VLAN is added.
   - If you are changing an existing routing VLAN, click the **Apply** button. Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

   The following table describes the view-only fields on the page.

Table 126. VLAN routing configuration information

Field	Description
Port	The interface for the routing VLAN.
MAC Address	The MAC address of the interface for the routing VLAN.

Remove the routing function from a VLAN

You can remove the routing function from a VLAN. That is, you can remove the routing IP address and subnet mask from the VLAN. The VLAN itself is not removed.

To remove the routing function from a VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > VLAN > VLAN Routing Configuration**.
   The VLAN Routing page displays.
3. Select the check box for the VLAN.
4. Click the **Delete** button.
   Your settings are saved. The routing function is removed from the VLAN but the VLAN itself is not deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Address Resolution Protocol

The Address Resolution Protocol (ARP) associates a Layer 2 MAC address with a Layer 3 IPv4 address. The switch support both a dynamic and manual ARP configuration. With manual ARP configuration, you can statically add entries into the ARP table.
ARP is a required part of the Internet Protocol (IP) and is used to translate an IP address to a media (MAC) address, defined by a LAN such as an Ethernet LAN. A device that sends IP packets must learn the MAC address of the IP destination, or of the next hop router, if the destination is not on the same subnet. This is achieved by broadcasting an ARP request packet, to which the intended recipient responds by unicasting an ARP reply that contains its MAC address. Once learned, the MAC address is used in the destination address field of the Layer 2 header prepended to the IP packet.
The ARP cache is a table maintained locally on each device on a network. ARP cache entries are learned by examining the source information in the ARP packet payload fields, regardless of whether it is an ARP request or response. Thus, when an ARP request is broadcast to all devices on a LAN or VLAN, each recipient can store the sender's IP and MAC address in its respective ARP cache. The ARP response (which is unicast message), is normally detected only by the requestor, who stores the sender information in its ARP cache. The most recent information always replaces existing content in the ARP cache.
The number of supported ARP entries depends on a device.
A device can be moved in a network, which means that the device's IP address that was associated with one MAC address is now associated with another MAC address. A device can also disappear from the network altogether (for example, it was reconfigured, disconnected, or powered off). These situations cause stale information in the ARP cache. Therefore, entries are updated or periodically refreshed to determine if an address still exists. If an entry was identified as a sender of an ARP packet, the entry can be removed from the ARP cache. You can configure an age-out interval that determines how long an entry that is not updated remains in the ARP

cache.

Display the ARP entries in the ARP cache

You can view ARP entries in the ARP cache. The ARP cache is a table that lists the remote connections that were recently detected by the switch

To display the ARP entries in the ARP cache:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > ARP> Basic > ARP Cache**.
   The ARP Cache page displays.
   The page displays the following information:
   - **IP Address**: Displays the IP address associated with the system's MAC address. This address must be the IP address of a device on a subnet attached to one of the switch's existing routing interfaces.
   - **Port**:Displays the associated unit/slot/port of the connection.
   - **MAC Address**: Displays the unicast MAC address of the device. The address is six 2-digit hexadecimal numbers separated by colons, for example, 00:06:29:32:81:40.

   The ARP cache be large. The pagination navigation menu functions as follows:
   - **Rows per page**: Select how many table entries are displayed per page. Possible values are 20, 50, 100, 200, and All. If you select All, the browser might be slow to display the information.
   - **<**. Display the previous page of the table data entries.
   - **>**. Display the next page of the table data entries.
3. To refresh the page, click the **Refresh** button.

Add or change a static entry in the ARP table

You can add a new static entry to the ARP table or change an existing entry that you manually added.

To add a new static entry or change an existing entry in the ARP table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > ARP > Advanced > ARP Create**.
   The ARP Create page displays.

3. If you are changing an existing ARP entry, select the check box for the entry.
4. In the **IP Address** field, specify the IP address to add.
   The address must be the IP address of a device on a subnet attached to one of the switch's existing routing interfaces.
5. In the **MAC Address** field, specify the unicast MAC address of the device.
   Enter the address as six 2-digit hexadecimal numbers separated by colons, for example, 00:06:29:32:81:40.
6. Do one of the following:
   - If you are adding a new ARP entry, click the **Add** button. Your settings are saved. The ARP entry is added.
   - If you are changing an existing ARP entry, click the **Apply** button. Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

   The ARP table can be large. The pagination navigation menu functions as follows:
   - **Rows per page**: Select how many table entries are displayed per page. Possible values are 20, 50, 100, 200, and All. If you select All, the browser might be slow to display the information.
   - **<**. Display the previous page of the table data entries.
   - **>**. Display the next page of the table data entries.

   The following table describes the view-only fields on the page.

Table 127.ARP cache information

Field	Description
IP Address	The IP address of the device.
Port	The associated interface (in the unit/port format) for the device connection.
MAC Address	The unicast MAC address of the device.
Type	The type of ARP entry. Possible values are as follows: • **Local**: An ARP entry associated with one of the switch's routing interface's MAC addresses. • **Gateway**: A dynamic ARP entry for which the IP address is that of a router. • **Static**: An ARP entry that you added. • **Dynamic**: An ARP entry that was learned by the switch.
Age	The time (in seconds) since the entry was last refreshed in the ARP table.

Delete a static ARP entry

You can delete a static ARP entry that you no longer need.

To delete a static ARP entry:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > ARP > Advanced > ARP Create**.
   The ARP Create page displays.
3. Select the check box for the entry.
4. Click the **Delete** button.
   Your settings are saved. The entry is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the ARP table settings or remove entries from the table

You can change the configuration settings for the ARP table. You can also remove entries from the table, for example, all dynamic entries.

To configure the ARP table settings or remove entries from the table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > ARP > Advanced > Global ARP Configuration**.
   The Global ARP Configuration page displays.
3. In the **Age Time** field, specify the period in seconds that a dynamic ARP entry remains in the ARP table before aging out.
   The range is from 15 to 21600 seconds. The default is 1200 seconds.
4. In the **Response Time** field, specify the period in seconds that the switch waits for an ARP response to an ARP request that it sends.
   The range is from 1 to 10 seconds. The default is 1 second.
5. In the **Retries** field, specify the maximum number of times an ARP request is retried after the switch does not receive an ARP response.
   This number includes the initial ARP request. The range is from 0 to 10. The default is 4.

6. In the **Cache Size** field, specify the maximum number of entries allowed in the ARP table. This number includes all static and dynamic ARP entries.The range is from 384 to 760. The default is 760.
7. Select the Dynamic Renew **Enable** or **Disable** radio button to enable or disable dynamic ARP entries to attempt to be automatically renewed after aging out.
   The default is Enable.
8. To remove certain entries from the ARP table, select one of the following options from the **Remove From Table** menu:
   - **All Dynamic Entries**: All dynamic entries will be removed.
   - **All Dynamic and Gateway Entries**: All dynamic and gateway entries will be removed.
   - **Specific Dynamic/Gateway Entry**: You must specify the IP address of the specific dynamic entry or specific gateway entry in the **Remove IP Address** field.
   - **Specific Static Entry**: You must specify the IP address of the specific static entry in the **Remove IP Address** field.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.
    The following table describes the view-only fields on the page.

Table 128.ARP Table Configuration

Field	Description
Total Entry Count	The total number of entries in the ARP table
Peak Total Entries	The highest value reached in the Total Entry Count field. This counter is restarted when you change the size of the ARP table cache.
Active Static Entries	The total number of active static entries in the ARP table
Configured Static Entries	The total number of configured static entries in the ARP table
Maximum Static Entries	The maximum number of static entries that can be defined

Routing Information Protocol

Routing Information Protocol (RIP) is a protocol that switches and routers can use to exchange network topology information. RIP is characterized as an interior gateway protocol, and is typically used in small to medium-sized networks. A device that runs RIP sends the contents of its routing table to each of its adjacent devices. When a route is removed from the routing table, it is flagged as unusable by the other devices and, after a certain period, removed from their routing.

Two versions of RIP exist:

- **RIPv1 (defined in RFC 1058)**:
  - Routes are specified by IP destination network and hop count
  - The routing table is broadcast to all devices on the attached network
- **RIPv2 (defined in RFC 1723)**:
  - The route specification also includes the subnet mask and gateway
  - The routing table is sent to a multicast address, reducing network traffic
  - Authentication is used for security

You can configure an interface to do the following:

- Receive packets in either or both formats
- Send packets formatted for RIPv1 or RIPv2, or send RIPv2 packets to the RIPv1 broadcast address
- Prevent any RIP packets from being received
- Prevent any RIP packets from being sent

Enable or disable RIP on the switch

By default, RIP is enabled on the switch.

To enable or disable RIP on the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > RIP > Basic > RIP Configuration**.
   The RIP Configuration page displays.
3. Select the RIP Admin Mode **Enable** or **Disable** radio button to enable or disable RIP on the switch.
   The default is Enable.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the global RIP settings for the switch

You can configure the global RIP settings for the switch, as opposed to the RIP settings for individual interfaces.

To configure global RIP settings for the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > RIP > Advanced > RIP Configuration**.
   The RIP Configuration page displays.
3. Select the RIP Admin Mode **Enable** or **Disable** radio button to enable or disable RIP on the switch.
   The default is Enable.
4. Select one of the following Split Horizon Mode radio buttons:
   - **None**: No special processing is enabled.
   - **Simple**: A route is not included in updates that are sent to the device from which it was learned. The default is Simple.
   - **Poison Reverse**: A route is included in updates sent to the device from which it was learned, but the metric is set to infinity.

   Split horizon is a technique for preventing problems that might be caused by including routes in updates sent to a device from which the route was originally learned.
5. Select the Auto Summary Mode **Enable** or **Disable** radio button to enable or disable groups of adjacent routes from being summarized in single entries.
   If you select the **Enable** radio button, adjacent routes are summarized in single entries to reduce the total number of entries. The default is Disable.
6. Select the Host Routes Accept Mode **Enable** or **Disable** radio button to enable or disable the switch from accepting host routes.
   If you select the **Enable** radio button, the switch accepts host routes. The default is Enable.
7. Select the Default Information Originate **Enable** or **Disable** radio button to enable or disable default route advertisements.
   The default is Disable.
8. In the **Default Metric** field, specify a metric value for redistributed routes.
   The range is from 1 to 15. The default is 0.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 129.RIP configuration information

Field	Description
Global Route Changes	The number of route changes made to the IP route database by RIP. This does not include the refreshing of a route after it aged out.
Global Queries	The number of responses sent to RIP queries from other devices.

Configure RIP interface settings

You can configure RIP settings that apply to individual interfaces.

To configure RIP interface settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > RIP > Advanced > Interface Configuration**.
   The Interface Configuration page displays.
3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

5. From the **Send Version** menu, select the version of RIP control packets that the interface sends:
   - **None**. No RIP control packets are sent.
   - **RIP-1**: Sends RIP version 1–formatted packets through broadcast.
   - **RIP-1c**: Send packets in RIP version 1–compatibility mode and sends RIP version 2–formatted packets through broadcast.
   - **RIP-2**: Send RIP version 2 packets using multicast. The default is RIP-2.
6. From the **Receive Version** menu, select the version of RIP control packets the interface accepts:
   - **RIP-1**: Accept RIP version 1–formatted packets only.
   - **RIP-2**: Accept RIP version 2–formatted packets only.
   - **Both**: Accept packets in either RIP version 1 or RIP version 2 format. The default is Both.
   - **None**: No RIP control packets are accepted.
7. From the **RIP Mode** menu, select **Enable** or **Disable** to enable or disable the RIP mode for the interface.
   Before you enable RIP version 1 or version 1c on an interface, you must first enable the network-directed broadcast mode (that is, the RIP mode) on the interface. The default is Disable.
8. From the **Authentication Type** menu, select the type of authentication:
   - **None**: No authentication protocols is configured. This is the default setting.
   - **Simple**: You are prompted to enter an authentication key. This key is included (in non-encrypted format) in the RIP header of all packets that are sent on the network. All devices on the network must be configured with the same key.
   - **Encrypt**: You are prompted to enter both an authentication key and an authentication ID. Encryption uses the MD5 Message-Digest algorithm. All devices on the network must be configured with the same key and ID.
9. If you select **Simple** or **Encrypt** from he Authentication Type menu, in the **Authentication Key** field, enter the RIP authentication key for the interface.
   The key can be up to 16 octets long.
10. If you select **Encrypt** from the Authentication Type menu, in the **Authentication Key ID** field, enter the RIP authentication Key ID for the interface.
    The key ID can from 0 to 255.
11. Click the **Apply** button.
    Your settings are saved.
12. To refresh the page, click the **Refresh** button.
13. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 130.RIP interface configuration information

Field	Description
Bad Packets Received	The number of RIP response packets received by the RIP process that were subsequently discarded for any reason
Bad Routes Received	The number of routes in valid RIP packets that were ignored for any reason (for example, unknown address family, or invalid metric)
Updates Sent	The number of triggered RIP updates sent on the interface. This explicitly does not include full updates sent containing new information.
IP Address	The IP address of the router interface
Link State	Indicates if the RIP router interface is up or down

Configure the RIP route redistribution settings and display the route redistribution summary

You can configure the RIP route redistribution settings and display the route redistribution summary.

To configure RIP route redistribution settings and display the route redistribution summary:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > RIP > Advanced > Route Redistribution**.
   The Route Redistribution page displays.
   The options that you can select from the **Source** menu consist of only those source routes that are already configured for redistribution by RIP. This allows you to configure another source route among the available source routes.
3. From the **Source** menu, select one of the following types of source routes:
   - **Connected**: Lets you configure the settings for the connected routes.
   - **Static**: Lets you configure the settings for the static routes.
   - **OSPF**: Lets you configure the settings for the OSPF routes.
4. From the **Redistribute Mode** menu, select **Enable** or **Disable** to enable or disable the RIP redistribute mode for the selected type of source route.
   The default is Disable.

5. In the **Metric** field, specify the metric value for the redistributed routes for the selected type of source route.
   The range is from 0 to 15. The default is 0, which means that no metric value is in effect.
6. In the **Distribute List** field, specify the access list that must filter the routes to be redistributed by the destination protocol.
   Only permitted routes are redistributed. If this field refers to a non-existent access list, all routes are permitted. The range for access list IDs is from 0 to 199. (For more information, see Access control lists (ACLs) on page 793.) The default is 0, which means that no access list is in effect.
   When you use an access list for route filtering, only the following fields in the access list are used:
   - Source IP address and netmask
   - Destination IP address and netmask
   - Action (permit or deny)

   All other fields (such as source and destination Port, precedence, ToS, and so on) are ignored. The source IP address is compared to the destination IP address of the route. The source IP netmask in the access list rule is treated as a wildcard mask, indicating which bits in the source IP address must match the destination address of the route.

   NOTE: In the ACL, if you enter 1 as a value for the netmask, the ACL ignores the corresponding address bit.

   When an access list rule includes a destination IP address and netmask (an extended access list), the destination IP address is compared to the network mask of the destination of the route. The destination netmask in the access list serves as a wildcard mask, indicating which bits in the route's destination mask are significant for the filtering operation.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 131.RIP route redistribution summary information

Field	Description
Source Protocol	The type of source route to be redistributed by RIP: • Connected • Static • OSPF
Redistribute Mode	Indicates if the route for a source protocol is redistributed. By default, route distribution is disabled.
Metric	The metric of redistributed routes for the source route. The field displays 0 if the metric is not configured.
Distribute List	The access list that filters the routes to be redistributed by the destination protocol. The field displays 0 if no ACL is configured for this purpose.
Match Internal	Indicates if an internal OSPF route is redistributed. By default, this option is enabled.
Match External Type 1	Indicates if an external type 1 OSPF route is redistributed. By default, this option is disabled.
Match External Type 2	Indicates if an external type 2 OSPF route is redistributed. By default, this option is disabled.
Match NSSA External Type 1	Indicates if an NSSA type 1 OSPF route is redistributed. By default, this option is disabled.
Match NSSA External Type 2	Indicates if an NSSA type 2 OSPF route is redistributed. By default, this option is disabled.

Router discovery and router advertisements

By default, a routing interface does not send router advertisements.You can enable the router advertisements for a routing interface and configure the setting for the router advertisements.

To configure router discovery for a routing interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Router Discovery > Router Discovery Configuration**.
   The Router Discovery Configuration page displays.

3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Advertise Mode** menu, select **Enable** or **Disable** to enable or disable the transmission of router advertisements on the interface.
   The default is Disable.
6. In the **Advertise Address** field, specify the IPv4 address that the interface must use for router advertisements.
7. In the **Maximum Advertise Interval** field, specify the maximum period in seconds that is allowed between router advertisements that are sent from the interface.
   The period must be in the range from 4 to 1800. Default value is 600.
8. In the **Minimum Advertise Interval** field, specify the minimum period in seconds that is allowed between router advertisements that are sent from the interface.
   The period must be in the range from 3 to 1800. The default is 450.
9. In the **Advertise Lifetime** field, specify the maximum period in seconds that the address that is advertised by the interface is considered a valid router address by hosts.
   The period must be in the range from 4 to 9000. Default is 1800.
10. In the **Preference Level** field, specify the preference level of the router interface as a default router relative to other routers on the same subnet.
    A higher number means that the advertized address receives a higher preference. The default is 0.

11. Click the **Apply** button.
    Your settings are saved.
12. To save the settings to the running configuration, click the **Save** icon.

Virtual Router Redundancy Protocol

If you configure a *static route* between a device and router (or switch) that is processing the device's routed traffic, you are effectively introducing a potential single point of failure into the network. If the router goes down, the device is unable to communicate. Virtual Router Redundancy Protocol (VRRP) can provide a backup mechanism through the use of a primary and secondary router.
If the primary router goes down, VRRP can let the secondary router take over from the primary router without affecting the device that uses the static route. The device uses a virtual IP address that is recognized by the secondary router if the primary router goes down.
You can configure physical ports or routed VLANs as VRRP router ports. The ports that participate in VRRP use an election protocol that determines which port functions as the primary router. You can configure multiple ports on the switch as virtual router ports. A single port can be visible to the network as more than one virtual router.

Enable VRRP and add a primary virtual router

You can enable VRRP and add a primary virtual router by configuring its virtual router ID (VRID), selecting its interface, and settings its primary IP address.

To enable VRRP and add a primary virtual router:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > VRRP > Basic > VRRP Configuration**.
   The VRRP Configuration page displays. The page shows different sections.
3. Select an Admin Mode radio button:
   - **Enable**: VRRP is enabled for the switch.
   - **Disable**: VRRP is disabled for the switch. This is the default setting.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

6. In the **VRID** field, enter a number in the range from 1 to 255.
   The VRID enables virtual routers to communicate with each other in a VRRP configuration. That is, the primary and secondary virtual routers must have the same VRID.
7. From the **Interface** menu, select the interface that must function as a primary virtual router.
8. In the **Primary IP Address** field, enter the IP address for the primary virtual router.
9. From the **Mode** menu, select if the virtual router is enabled:
   - **Active**: The primary virtual router is enabled.
   - **Inactive**: The primary virtual router is configured but disabled.
10. Click the **Add** button.
    Your settings are saved. The primary virtual router is added.
11. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
Interface IP Address	The IP address that is associated with the interface. This IP address can be different from the primary IP address.
State	The state of the virtual router: • **Initializing**: The virtual router is in the process of initializing. • **Master**: The virtual router completed initialization and functions as the primary router. • **Backup**: The virtual router completed initialization and functions as the secondary router.

Enable VRRP and add a primary virtual router with enhanced settings

You can enable VRRP and add a primary virtual router by configuring its virtual router ID (VRID), selecting its interface, and settings its primary IP address. In addition, you can configure multiple enhanced settings.

To enable VRRP and add a primary virtual router with enhanced settings

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > VRRP > Advanced > VRRP Configuration**.
   The VRRP Configuration page displays. The page shows different sections.

3. Select an Admin Mode radio button:
   - **Enable**: VRRP is enabled for the switch.
   - **Disable**: VRRP is disabled for the switch. This is the default setting.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.
6. In the **VRID** field, enter a number in the range from 1 to 255.
   The VRID enables virtual routers to communicate with each other in a VRRP configuration. That is, the primary and secondary virtual routers must have the same VRID.
7. From the **Interface** menu, select the interface that must function as the primary virtual router.
8. From the **Pre-empt Mode** menu, select if a backup virtual router can preempt the primary virtual router:
   - **Enable**: A backup VRRP router can preempt the primary virtual router if the backup router has a greater priority than the primary virtual router *and* the IP address in the Primary IP Address field (see below) is not the IP address of the primary virtual router. The default is Enable.
   - **Disable**: A backup virtual router cannot preempt the primary virtual router.
9. From the **Accept Mode** menu, select if the primary virtual router can accept data packets:
   - **Enable**: The primary virtual router accepts all types of data packets that are addressed to the IP address that is associated with the primary virtual router.
   - **Disable**: The primary virtual router discards all types of data packets that are addressed to the IP address that is associated with the primary virtual router, unless the IP address in the Primary IP Address field (see below) *is* the IP address of the primary virtual router. The default is Disable.
10. In the **Configured Priority** field, enter the priority number that VRRP uses in the election of the primary virtual router.
    The range is from 1 to 254. If the virtual IP address is the same as the interface IP address, the priority is always set to 254.
11. In the **Advertisement Interval** field, enter the time, in seconds, between the transmission of advertisement packets by the router.
    Enter a number from 1 to 255. The default value is 1 second.
12. In the **Primary IP Address** field, enter the IP address for the primary virtual router.

13. From the **Authentication Type** menu, select if authentication is enabled for the primary virtual router.
    - **None**: No authentication is performed. This is the default.
    - **Simple**: Authentication is performed using a text password.
14. If you select **Simple** from the **Authentication Type** menu, enter a password in the **Authentication Data** field.
15. From the **Status** menu, select if the primary virtual router is enabled:
    - **Active**: The primary virtual router is enabled.
    - **Inactive**: The primary virtual router is configured but disabled.
16. Click the **Add** button.
    Your settings are saved. The primary virtual router is added.
17. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
Operational Priority	The priority that is used for the VRRP election process of the primary virtual router. A higher value is a higher priority. A primary virtual router is assigned a priority of 0 if it no longer participates in VRRP because a backup virtual router took over. A virtual router is assigned a priority of 255 if the interface IP address is identical to the primary IP address.
Interface IP Address	The IP address that is associated with the interface. This IP address can be different from the primary IP address.
Owner	Indicates if the interface IP address is identical to the primary IP address: • **True**: The interface IP address is identical to the primary IP address. In a VRRP election process, if the virtual router is active and its interface IP address is identical to the primary IP address, it always wins an election for the primary virtual router. • **False**: The interface IP address is different from the primary IP address.
VMAC Address	The virtual MAC address that is associated with the virtual router. This MAC address consists of a 24-bit organizationally unique identifier, the 16-bit constant identifying the VRRP address block. and the 8-bit VRID.
State	The state of the virtual router: • **Initializing**: The virtual router is in the process of initializing. • **Master**: The virtual router completed initialization and functions as the primary router. • **Backup**: The virtual router completed initialization and functions as the secondary router.

Change the settings for a primary virtual router

You can change the settings for an existing primary virtual router.

To change the settings for a primary virtual router:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > VRRP > Advanced > VRRP Configuration**.
   The VRRP Configuration page displays. The page shows different sections.
3. Select the check box for the VRRP ID (VRID), which identifies the primary virtual router.

4. Change the settings as needed.
   For more information about the settings, see Enable VRRP and add a primary virtual router with enhanced settings on page 396.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Remove a primary virtual router

You can remove a primary virtual router that you no longer need.

To remove a primary virtual router:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > VRRP > Advanced > VRRP Configuration**.
   The VRRP Configuration page displays. The page shows different sections.
3. Select the check box for the VRRP ID (VRID), which identifies the primary virtual router.
4. Click the **Delete** button.
   Your settings are saved. The primary virtual router is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure a secondary virtual router

A virtual router that is assigned a VRRP secondary IP address can become the new primary virtual router if the original virtual router goes down. The primary and secondary virtual routers must have the same virtual router ID (VRID).

To configure a secondary virtual router:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > VRRP > Advanced > VRRP Secondary IP Address Configuration**.
   The VRRP Secondary IP Address Configuration page displays. The page shows different sections.
3. From the **VRRP Interface** menu, select the VRRP interface that must function as a secondary virtual router.

4. From the **VRRP ID** menu, select the VRID that is assigned to the secondary virtual router.
5. In the **Secondary IP Address** field, enter the IP address for the secondary virtual router.
   This address must be part of an IP address subnet on the interface.
   The Primary IP Address fields shows the IP address that is assigned to the primary virtual router with the same VRID.
6. Click the **Add** button.
   Your settings are saved. The secondary virtual router is added.
7. To save the settings to the running configuration, click the **Save** icon.

Change the IP address for a secondary virtual router

You can change the IP address for an existing secondary virtual router.

To change the IP address for a secondary virtual router:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > VRRP > Advanced > VRRP Secondary IP Address Configuration**.
   The VRRP Secondary IP Address Configuration page displays. The page shows different sections.
3. From the **VRRP Interface** menu, select the VRRP interface.
4. From the **VRRP ID** menu, select the VRID.
5. Select the check box for the secondary IP address.
6. In the **Secondary IP Address** field, change the IP address.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Remove a secondary virtual router

You can remove a secondary virtual router that you no longer need.

To remove a secondary virtual router:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Routing > VRRP > Advanced > VRRP Secondary IP Address Configuration**.
   The VRRP Secondary IP Address Configuration page displays. The page shows different sections.
3. From the **VRRP Interface** menu, select the VRRP interface.
4. From the **VRRP ID** menu, select the VRID.
5. Select the check box for the secondary IP address.
6. Click the **Delete** button.
   Your settings are saved. The secondary virtual router is removed.
7. To save the settings to the running configuration, click the **Save** icon.

Configure VRRP interface tracking and route tracking

You can let VRRP track specific interfaces and routes and let VRRP determine if they can be reached. If a tracked interface or route is not reachable, VRRP can let a secondary router take over.

To configure VRRP interface tracking and route tracking:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > VRRP > Advanced > VRRP Tracking Configuration**.
   The VRRP Tracking Configuration page displays. The page shows different sections.
3. From the **VRRP Interface** menu, select the VRRP interface that you want to track.
4. From the **VRRP ID** menu, select the virtual router ID (VRID) that you want to track.
5. In the VRRP Tracking Interface Configuration section, configure the following settings:
   - From the **Tracked Interface** menu, select the routing interface to be tracked. The menu shows all routing interfaces that are not yet tracked for the selected VRRP interface and VRRP ID combination. Loopback interfaces and tunnels that cannot be tracked are not shown.
   - In the **Priority Decrement** field, enter a priority decrement value for the tracked interface. If the interface cannot be reached, its VRRP priority is decreased by the value that you enter. The range is from 1 to 254. The default is 10.

   The Tracked Interface State field displays if the interface can be reached.

6. In VRRP Tracking Route Configuration section, configure the following settings:
   - In the **Tracked Route Prefix** field, enter the prefix of the route to be tracked.
   - In the **Tracked Route Prefix Length** field, enter the prefix length of the route to be tracked.
   - In the **Priority Decrement** field, enter a priority decrement value for the tracked route. If the route cannot be reached, its VRRP priority is decreased by the value that you enter. The range is from 1 to 254. The default is 10.

   The Reachable field displays if the tracked route can be reached.
7. Click the **Add** button.
   Your settings are saved. The tracking configuration is added.
8. To save the settings to the running configuration, click the **Save** icon.

Display VRRP statistics

If VRRP is enabled, you can display global VRRP statistics and statistics for each VRRP ID.

To display VRRP statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Advanced > VRRP Statistics**.
   The VRRP Statistics page displays. The page shows different sections.
3. To refresh the page, click the **Refresh** button.

The following table describes the view-only fields on the page.

Field	Description
Global Statistics	
Router Checksum Errors	The number of VRRP packets received with an invalid VRRP checksum value
Router Version Errors	The number of VRRP packets received with an unknown or unsupported version number
Router VRID Errors	The number of VRRP packets received with an invalid VRID
Statistics	
VRRP ID	The VRRP ID (VRID)
Interface	The interface that is configured as a virtual router
Up Time	The time, in days, hours, minutes and seconds, that elapsed since the virtual router transitioned to the initialized state
State Transitioned to Master	The number of times that the virtual router's state transitioned to primary
Advertisement Received	The number of VRRP advertisements received by the virtual router
Advertisement Interval Errors	The number of VRRP advertisement packets received for which the advertisement interval was different from the one configured for the virtual router
Authentication Failure	The number of VRRP packets received that did not pass the authentication check
IP TTL Errors	The number of VRRP packets received by the virtual router with a time-to-live (TTL) value that is not equal to 255
Zero Priority Packets Received	The number of VRRP packets received by the virtual router with a priority of 0
Zero Priority Packets Sent	The number of VRRP packets sent by the virtual router with a priority of 0
Invalid Type Packets Received	The number of VRRP packets received by the virtual router with an invalid type
Address List Errors	The number of packets received for which the address list does not match the locally configured list for the virtual router
Invalid	The number of packets received with an unknown authentication type

Field	Description
Authentication Type	
Authentication Type Mismatch	The number of packets received with an authentication type different from the locally configured authentication method
Packet Length Errors	The total number of packets received with a packet length that is smaller than the length of the VRRP header

Configure OSPF and OSPFv3

This chapter covers the following topics:

- Open Shortest Path First
- Open Shortest Path First version 3

NOTE: In this chapter, open shortes path first (OSPF) refers to OSPF version 2 (OSPFv2). OSPFv3 refers to OSPF version 3.

Open Shortest Path First

Open Shortest Path First (OSPF) lets the switch determine the best (shortest) route to a destination. For large networks, OSPF is generally used in preference to Routing Information Protocol (RIP).

If you manage a large or complex network, OSPF offers several benefits:

- **Less network traffic**:
  - Routing table updates are sent only when a change occurs
  - Only the part of the table that changes is sent
  - Updates are sent to a multicast address rather than a broadcast address
- **Hierarchical management**: Hierarchical management allows the network to be subdivided. The top level of the hierarchy of an OSPF network is known as an autonomous system (AS) or routing domain. This is a collection of networks with a common administration and routing strategy. The AS is divided into areas:
  - **Intra-area**: Intra-area routing is used when a source and destination address are in the same area.
  - **Inter-area**: Inter-area routing across an OSPF backbone is used when a source and destination address are *not* in the same area. An inter-area router communicates with border routers in each of the areas to which it provides connectivity.

  If you enable OSPF on the switch, it determines the best route using the assigned cost and the type of the OSPF route. If more than one type of route is present, the switch selects a route based on the following order of preference:
  1. Intra-area route
  2. Inter-area route
  3. External type 1 route: The route is external to the AS
  4. External type 2 route: The route is learned from other protocols such as RIP

Enable OSPF

Before OSPF can become operational, you must enable the administrative mode for IPv4 routing on the switch (see Configure the Routing IP Settings on page 307 on page 354).

To enable OSPF:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Basic > OSPF Configuration**.
   The OSPF Configuration page displays.
3. Select an Admin Mode radio button:
   - **Enable**: OSPF is enabled for the switch. This is the default setting.
   - **Disable**: OSPF is disabled for the switch.

   The Router ID field displays the 32-bit number in dotted-decimal format. This number uniquely identifies the switch within the autonomous system (AS).
   To change the router ID, you must first disable OSPF. After you set a new router ID, you must reenable OSPF for the change to take effect. The default value is 0.0.0.0, although this is not a valid router ID.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the OSPF default route advertisement

Whether or not a default route exists in the routing table, you can set up a default route for OSPF advertisements, including a metric value and a type of metric.

To configure the OSPF default route advertisement:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > OSPF Configuration**.
   The OSPF Configuration page displays. The page shows different sections.
3. Go to the Default Route Advertise Configuration section.

4. Select a Default Information Originate radio button:
   - **Enable**: OSPF originates a default route for external LSA advertising. This route is in the format X.X.X.X/X.X.X.X.
   - **Disable**: OSPF does not originate a default route for external LSA advertising. This is the default settings.
5. Select an Always radio button:
   - **True**: OSPF originates a default route regardless of whether a default route exists.
   - **False**: OSPF originates a default route only if a default route is already in the switch's routing table.
6. In the **Metric** field, set the metric value of the default route.
   The value ranges from 0 to 16777214. The default is 0.
7. Select a Metric Type field radio button:
   - **External Type 1**: This metric is computed by adding the internal OSPF cost to the external redistributed cost.
   - **External Type 2**: This metric is computed by using only the external redistributed cost. This is the default.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Configure the global OSPF settings

You can configure the OSPF settings that apply globally to the OSPF processes on the switch.

To configure the global OSPF settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > OSPF Configuration**.
   The OSPF Configuration page displays. The page shows different sections.
3. Go to the OSPF Configuration section.
4. In the **Router ID** field, enter the 32-bit number in dotted-decimal format that uniquely identifies the router within the autonomous system (AS).
   To change the router ID, you must first disable OSPF. After you set a new router ID, you must reenable OSPF for the change to take effect. The default value is 0.0.0.0, although this is not a valid router ID.

5. Select an Admin Mode radio button:
   - **Enable**: OSPF is enabled for the switch. This is the default setting.
   - **Disable**: OSPF is disabled for the switch.

   Before OSPF can become operational, you must enable the administrative mode for IPv4 routing on the switch (see Manage the global IPv4 routing settings on page 354).
6. From the **RFC 1583 Compatibility** menu, select how the switch automatically sets preference rules for link-state advertisements (LSAs) if multiple AS-external-LSAs advertise the same destination:
   - **Enable**: The preference rules prevent routing loops when AS-external-LSAs for the same destination originate from different areas. All routers or switches in the OSPF domain must be configured the same way. Enable is the default setting. For more information, see section 16.4.1 of the OSPF-2 standard (RFC 2328): https://datatracker.ietf.org/doc/html/rfc2328
   - **Disable**: If all OSPF routers are capable of operating according to the standard described in RFC 2328, disable RFC 1583 Compatibility.
7. From the **Opaque LSA Status** menu, select how the switch treats opaque LSAs:
   - **Enable**: The switch stores opaque LSAs. An opaque LSA is used for flooding user-defined information within an OSPF router domain.
   - **Disable**: The switch does not store opaque LSAs.
8. In the **Exit Overflow Interval** field, set how many seconds OSPF must wait before leaving an overflow state.
   When the number of nondefault external LSAs exceeds a limit, the switch enters an overflow state in which OSPF cannot originate nondefault external LSAs. (For more information, see RFC 1765.)
   If the number of seconds is 0 (the default setting), OSPF does not leave the overflow state until you disable and reenable OSPF. The maximum possible value is 2,147,483,647 seconds.
9. In the **SPF Delay Time** field, set how many seconds OSPF must wait before starting a new Shortest Path First (SPF) calculation.
   The SPF delay time is the number of seconds from when OSPF receives a topology change to the start of the next SPF calculation. The delay can be from 0 to 65535 seconds. The default is 5 seconds. A value of 0 means that there is no delay: the SPF calculation is started as soon as a topology change occurs.
10. In the **SPF Hold Time** field, set how many seconds OSPF must wait before between two consecutive SPF calculations.
    The SPF hold time is the minimum time in seconds between two consecutive SPF calculations. The hold time can be from 0 to 65,535 seconds. The default is 10 seconds. A value of 0 means that there is no delay: two SPF calculations can be done, one immediately after the other.

11. In the **External LSDB Limit** field, set the limit of the external link-state database (LSDB) for OSPF.
    When the number of nondefault AS-external-LSAs in a router's link state database reaches the external LSDB's limit, the router enters the overflow state. The router cannot hold more none-default AS-external-LSAs than the maximum number that the external LSDB can hold. The external LSDB's limit must be set identically in all routers that are attached to the OSPF backbone or any regular OSPF area. The limit can be from –1 to 2147483647. The default setting is –1, which indicates that there is no limit.
12. In the **Default Metric** field, set the default for the metric of redistributed routes.
    The range is 1 from 16777214. The default is 0.
13. In the **Maximum Paths** field, set the number of paths that OSPF can report for a destination.
    The range is from 1 to 16. The default is 4.
14. In the **AutoCost Reference Bandwidth** field, set the bandwidth that controls how OSPF calculates link cost.
    Set the reference bandwidth in megabits per second. Unless a link cost is configured, the link cost is computed by dividing the reference bandwidth by the interface bandwidth. The range is from 1 to 4294967. The default is 100.
15. From the **Default Passive Setting** menu, select how the global passive mode setting for all OSPF interfaces is configured:
    - **Enable**: Enabling the passive mode globally overwrites any passive mode setting at the interface level. OSPF does not form adjacencies on passive interfaces, but does advertise attached networks as stub networks.
    - **Disable**: The default is Disabled.
16. Click the **Apply** button.
    Your settings are saved.
17. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
ASBR Mode	The router is an autonomous system boundary router (ASBR) if it is configured to redistribute routes from another protocol, or if it is configured to originate an external LSA advertising the default route.
ABR Status	The router is an area border router (ABR) if it has active non-virtual interfaces in two or more OSPF areas.
External LSA Count	The number of external (LS type 5) link state advertisements (LSAs) in the link state database.
External LSA Checksum	The sum of the LS checksums of the external LSAs that are contained in the link state database. You can use this sum to determine if there was a change in the link state database of a switch or router, and to compare the link state databases of two devices. This value is in hexadecimal.
AS_OPAQUE LSA Count	The number of opaque LSAs with domain-wide flooding scope.
AS_OPAQUE LSA Checksum	The sum of the LS checksums of the opaque LSAs with a domain-wide flooding scope. You can use this sum to determine if there was a change in the link state database of a switch or router, and to compare the link state databases of two devices.
New LSAs Originated	The number of LSAs originated by the switch. In any OSPF area, a router (or in this situation, a switch) originates several LSAs. Each router originates a router-LSA. If the router is also the designated router for any of the area's networks, it originates network LSAs for those networks.
LSAs Received	The number of received LSAs for new instantiations. This number does not include new instantiations of self-originated LSAs.

Add or delete an OSPF common area ID

An OSPF common area consists of a group of devices that share the same area ID.

To add or delete an OSPF common area ID:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > Common Area Configuration**.
   The OSPF Common Area Configuration page displays.

3. In the **Area ID** field, enter an OSPF area ID.
   An area ID is a 32-bit integer in dotted-decimal format that identifies the area to which a switch interface connects.
4. Take one of the following actions:
   - **Add**: Click the **Add** button. The area ID is added.
   - **Delete**: Click the **Delete** button. The area ID is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
External Routing	A definition of the switch's capabilities for the area, including whether or not AS-external-LSAs are flooded into or throughout the area. If the area is a stub area, one of the following external routing capabilities can display: • **Import External LSAs**: Imports and propagates external LSAs. • **Import No LSAs**: Does not import and propagate external LSAs. If the area is *not* a stub area, the only option is Import No LSAs.
SPF Runs	The number of times that the intra-area route table was calculated using this area's link state database. (This is typically calculated using Dijkstra's algorithm.)
Area Border Router Count	The total number of area border routers that re reachable within this area. This is initially zero, and is calculated in each SPF pass.
Area LSA Count	The total number of link state advertisements in this area's link state database, excluding AS external LSAs.
Area LSA Checksum	The 32-bit unsigned sum of the link state advertisement checksum contained in this area's link state database. This sum excludes external (LSA type 5) link state advertisements. You can use the sum to determine if a change occurred in the link state database of the switch, and to compare the link state database of two devices.
Flood List Length	The number of LSAs on this area's flood list.
Import Summary LSAs	The summary LSAs that are imported into this area.

Add an OSPF stub area

A stub area is a controlled area that does not allow external routes from outside the OSPF network.

To add an OSPF stub area:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > Stub Area Configuration**.
   The OSPF Stub Area Configuration page displays.
3. In the **Area ID** field, enter an OSPF area ID.
   An area ID is a 32-bit integer in dotted-decimal format that identifies the area to which a switch interface connects.
4. From the **Import Summary LSAs** menu, select if summary LSAs are imported:
   - **Enable**: Summary LSAs are imported into stub areas.
   - **Disable**: Summary LSAs are not imported into stub areas.
5. In the **Default Cost** field, enter the metric value that must be applied for the default route advertised to the stub area.
   The range is from 1 to 16,777,215.
6. Click the **Add** button.
   The area is added as a stub area.
7. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
SPF Runs	The number of times that the intra-area route table was calculated using this area's link state database. (This is typically calculated using Dijkstra's algorithm.)
Area Border Router Count	The total number of area border routers that are reachable within this area. This is initially zero, and is calculated in each SPF pass.
Area LSA Count	The total number of link state advertisements in this area's link state database, excluding AS external LSAs.
Area LSA Checksum	The 32-bit unsigned sum of the link state advertisement checksum contained in this area's link state database. This sum excludes external (LSA type 5) link state advertisements. You can use the sum to determine if a change occurred in the link state database of the switch, and to compare the link state database of two devices.
Type of Service	The normal ToS that is associated with the stub metric.

Add an OSPF NSSA area

The OSPF not-so-stubby area (NSSA) can receive external routes (Type 7 LSAs) and learn routes from other protocols such as RIP and forward these routes to other areas. You can select if summary LSAs (Type 3 LSAs) can be imported into stub areas.

To add an OSPF NSSA area:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > NSSA Area Configuration**.
   The OSPF NSSA Area Configuration page displays.
3. In the **Area ID** field, enter an OSPF area ID.
   An area ID is a 32-bit integer in dotted-decimal format that identifies the area to which a switch interface connects.
4. From the **Import Summary LSAs** menu, select if summary LSAs are imported:
   - **Enable**: Summary LSAs are imported into stub areas.
   - **Disable**: Summary LSAs are not imported into stub areas.
5. The options in the Default Information Originate section of the table let you configure the switch to advertise a default route into the NSSA when the import of summary LSAs is disabled:
   a. From the **Admin Mode** menu, select **Enable** to let the switch advertise a default route or **Disable** to prevent the switch from doing so.
   b. In the **Metric Value** field, set the metric value for the default route. The range is from 1 to 16777214.
   c. From the **Metric Type** menu, select the type of metric that must be used in the default route:
      - **Comparable Cost**: External type 1 metrics that are comparable to the OSPF metric
      - **Non-comparable Cost**: External type 2 metrics that are assumed to be larger than the cost of the OSPF metric.
6. From the **Translator Role** menu, select the role of the NSSA:
   - **Always**: The switch assumes the role of the translator when it becomes a border router.
   - **Candidate**: The switch participates in the translator election process when it attains border router status.
7. In the **Translator Stability Interval** field, enter the period during which the switch extends its translator role after another device becomes the translator.
   The range is from 0 to 3600 minutes.

8. From the **Redistribute Mode** menu, select if learned external routes are redistributed:
   - **Enable**: The NSSA ABR distributes learned external routes to the NSSA.
   - **Disable**: The NSSA ABR does not distribute learned external routes to the NSSA.
9. Click the **Add** button.
   The area is added as an NSSA area.
10. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
SPF Runs	The number of times that the intra-area route table was calculated using this area's link state database. (This is typically calculated using Dijkstra's algorithm.)
Area Border Router Count	The total number of area border routers that are reachable within this area. This is initially zero and is calculated in each SPF pass.
Area LSA Count	The total number of link state advertisements in this area's link state database, excluding AS external LSAs.
Area LSA Checksum	The 32-bit unsigned sum of the link state advertisement checksum contained in this area's link state database. This sum excludes external (Type 5) link state advertisements. You can use the sum to determine if a change occurred in the link state database of the switch, and to compare the link state database of two devices.
Translator State	This field displays if and how the NSSA border router translates Type 7 LSAs into Type 5 LSAs: • **Enabled**: The NSSA border router's translator role is set to always. • **Elected**: The candidate NSSA border router is translating Type 7 LSAs into Type 5 LSAs. • **Disabled**: The candidate NSSA border router is not translating Type 7 LSAs into Type 5 LSAs.

Add an OSPF area range

You add an OSPF area range by setting an area ID with an address range (an IP address with a subnet mask). You then also set the LSAs that are associated with the area and whether the address range is advertised outside the area.
You can associate an OSPF area range with a common area (see Add or delete an OSPF common area ID on page 411), stub area (see Add an OSPF stub area on page 412), or NSSA area (see Add an OSPF NSSA area on page 414).

To add an OSPF area range:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > Area Range Configuration**.
   The OSPF Area Range Configuration page displays.
3. In the **Area ID** field, enter an OSPF area ID.
   An area ID is a 32-bit integer in dotted-decimal format that identifies the area to which a switch interface connects.
4. In the **IP Address** field, enter the IP address for the address range.
5. In the **Subnet Mask** field, enter the subnet mask for the address range.
6. From the **LSDB Type** menu, select the type of link state advertisements (LSAs) that must be associated with the area ID and address range.
   - **Network Summary**: The switch generates Network Summary LSAs (Type 3 LSAs) for the area. This is the default setting.
   - **NSSA External**: The switch generates NSSA External LSAs (Type 7 LSAs) for the area. These types of LSAs provide more information.
7. From the **Advertise** menu, select if the address range is advertised outside the area:
   - **Enable**: The address range is advertised outside the area through Network Summary LSAs. This is the default setting.
   - **Disable**: The address range is not advertised outside the area.
8. Click the **Add** button.
   The OSPF ID and area range are added.
9. To save the settings to the running configuration, click the **Save** icon.

Configure an OSPF interface

You can configure one or more OSPF interfaces, which are interfaces that are connected to OSPF areas over point-to-point or broadcast connections.

To configure an OSPF interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > Interface Configuration**.
   The OSPF Interface Configuration page displays.
3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces or VLANs by taking one of the following actions:
   - To configure a single interface or VLAN, select the check box associated with the interface or VLAN, or type the interface or VLAN number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces or VLANs with the same settings, select the check box associated with each interface or VLAN.
5. In the **Area ID** field, enter an OSPF area ID.
   An area ID is a 32-bit integer in dotted-decimal format that identifies the area to which a switch interface connects.

6. Select an Admin Mode radio button:
   - **Enable**: OSPF is enabled for the interface.
   - **Disable**: OSPF is disabled for the interface. This is the default setting.

   OSPF can become operational on the interface only if the interface is configured with an IP address and subnet mask (see Configure IPv4 routing interfaces on page 359).
   You can configure the OSPF settings without enabling OSPF on the interface, but the change does not take effect until you enable OSPF on the interface. The following fields display only if you enable OSPF on the interface:
   State, Designated Router, Backup Designated Router, Number of Link Events, Local Link LSAs, and Local Link LSA Checksum.
7. In the **Router Priority** field, enter the OSPF priority for the interface.
   This is a number from 0 to 255. The default is 1, which is the highest router priority. A value of 0 indicates that the router is not eligible to become the designated router on this network.
8. In the **Retransmit Interval** field, enter the number of seconds between link-state advertisements for adjacencies, that is, for neighboring devices that form an OSPF adjacency with the interface.
   The interval is also used when retransmitting database descriptions and link-state request packets. The value for the interval ranges from 1 to 3600 seconds (1 hour). The default is 5 seconds.
9. In the **Hello Interval** field, enter the OSPF hello interval for the interface.
   This interval must be the same for all routers and switches in the network. The interval ranges from 1 to 65,535 seconds. The default is 10 seconds.
10. In the **Dead Interval** field, enter the number of seconds that the interface waits for a neighbor's hello packets before determining that the neighbor is down.
    This interval must be the same for all routers and switches in the network. This interval must be a multiple of the hello interval. The interval ranges from 1 to 65,535. The default is 40 seconds. (The default hello interval is 10 seconds.)
11. In the **Iftransit Delay Interval** field, enter the estimated number of seconds required to transmit a link state update packet over the interface.
    The interval range from 1 to 3600 seconds (1 hour). The default is 1 second.
12. From the In the **MTU Ignore** menu, select if OSPF MTU mismatch detection is disabled:
    - **Enable**: OSPF MTU mismatch detection is *disabled* on incoming database description packets. The default value is Disable (MTU mismatch detection is enabled).
    - **Disable**: OSPF MTU mismatch detection is *enabled* on incoming database description packets. This is the default setting.

13. From the **Passive Mode** menu, select if the interface does not form an adjacency:
    - **Enable**: The interface is passive to prevent OSPF from forming an adjacency. OSPF advertises networks attached to passive interfaces as stub networks.
    - **Disable**: The interface is not passive and OSPF can form an adjacency. This is the default setting.
14. From the **Network Type** menu, select the type of OSPF network for the interface:
    - **Broadcast**: The interface can connect to multiple OSPF routers. This is the default setting.
    - **Point to point**: The interface can connect to one other OSPF router.
15. From the **Authentication Type** menu, select if communication over the interface is authenticated and encrypted:
    - **None**: The interface does not use authentication and encryption. This is the default settings.
    - **Simple**: The interface prompts you to enter an authentication key. This key is included (non-encrypted) in the OSPF header of all packets sent on the network. All routers and switches on the network must be configured with the same key.
    - **Encrypt**: The interface prompts you to enter an authentication key and an authentication ID. Encryption uses the MD5 Message-Digest algorithm. All routers and switches on the network must be configured with the same key and ID.
16. In the **Authentication Key** field, enter the OSPF authentication key for the interface.
    If you select **Simple** from the **Authentication Type** menu, you can set a key with maximum of 8 octets. If you select **Encrypt** from the **Authentication Type** menu, you can set a key with maximum of 16 octets.
17. In the **Authentication Key ID** field, enter the ID that must be used for authentication on the interface.
    If you select **Encrypt** from the **Authentication Type** menu, you must set an ID, which is a number between 0 and 255.
18. In the **Metric Cost** field, enter the value that OSPF uses to compute the shortest paths.
    The range is from 1 to 65,535. The default is 1.
19. Click the **Apply** button.
    Your settings are saved.
20. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
IP Address	The IP address of the interface
Subnet Mask	The network mask, indicating the portion of the IP address that identifies the attached network
LSA Ack Interval	The number of seconds that the interface must wait before sending a delayed acknowledgement
State The state of the OSPF interface:	
Down: This is the initial interface state. The interface is not usable for OSPF, and the interface settings are at their initial values. All interface timers are disabled, and no adjacencies are associated with the interface.	
Loopback: The interface is looped back either in hardware or software. The interface is unavailable for regular data traffic. You can get information about the interface by sending ICMP pings to the interface or through a bit error test because IP packets can still be addressed to an interface in a loopback state. An interface in a loopback state is advertised in LSAs as single host route for which the destination is the IP interface address.	
Waiting: The switch is monitoring received hello packets to detect the identity of the backup designated router for the network. The switch cannot elect a backup designated router or a designated router until it transitions out of the waiting state. This prevents unnecessary changes of the backup designated router.	
Designated Router: The switch is the designated router on the attached network. Adjacencies are established to all other routers and switches that in the network. The switch must also originate a network LSA for the network node. The network LSA contains links to all switches and routers (including the designated router) in the network.	
Backup Designated Router: The switch is the backup designated router on the attached network. The switch is promoted to designated router if the current designated router fails. The switch establishes adjacencies to all other routers attached to the network. The backup designated router performs slightly different functions during the LSA flooding, as compared to the designated router.	
Other Designated Router: The interface is connected to a broadcast on which other routers are the designated router and backup designated router. The switch attempts to form adjacencies to both the designated router and the backup designated router.	
Designated Router	The identity of the designated router for this network, as detected by the advertising router. The designated router is identified by its router ID. The value 0.0.0.0 means that no designated router is detected.

Field	Description
Backup Designated Router	The identity of the backup designated router for this network, as detected by the advertising router. The backup designated router is identified by its router ID. The value 0.0.0.0 means that no backup designated router is detected.
Number of Link Events	The number of times the OSPF interface changed its state
Local Link LSAs	The number of opaque LSAs for which the flooding scope is the link on the interface
Local Link LSA Checksum	The sum of the checksums of local link LSAs for the link

View or clear OSPF statistics for an interface

If OSPF is enabled, you can view or clear OSPF statistics for an interface.

To view or clear OSPF statistics for an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > Interface Statistics**.
   The OSPF Interface Statistics page displays.
3. From the **Interface** menu, select the interface.
4. To refresh the page, click the **Refresh** button.
5. To clear all the counters, resetting all OSPF interface statistics to default values, click the **Clear** button.
6. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
OSPF Area ID	The OSPF area to which the interface belongs. An OSPF area ID is a 32-bit integer in dotted-decimal format that uniquely identifies the area to which the interface connects.
Area Border Router Count	The number of area border routers (ABR) reachable within this area. This is initially zero, and is calculated in each Shortest Path First (SPF) pass.
AS Border Router Count	The total number of autonomous system (AS) border routers reachable within this area. This is initially zero, and is calculated in each SPF pass.
Area LSA Count	The number of link state advertisements in this area's link state database, excluding AS external LSAs.
IP Address	The IP address of the interface.
Interface Events	The number of times the OSPF interface changed its state, or an error occurred.
Virtual Events	The number of state changes or errors that occurred on this virtual link.
Neighbor Events	The number of times this neighbor relationship changed state, or an error occurred.
Sent Packets	The number of OSPF packets transmitted on the interface.
Received Packets	The number of valid OSPF packets received on the interface.
Discards	The number of received OSPF packets discarded because of an error in the packet or an error in processing the packet.
Bad Version	The number of received OSPF packets for which the version field in the OSPF header does not match the version of the OSPF process handling the packet.
Source Not On Local Subnet	The number of received packets discarded because the source IP address is not within a subnet configured on a local interface.
Virtual Link Not Found	The number of received OSPF packets discarded because the ingress interface is in a non-backbone area and the OSPF header identified the packet as belonging to the backbone, but OSPF does not have a virtual link to the packet's sender.
Area Mismatch	The number of OSPF packets discarded because the area ID in the OSPF header is not the area ID configured on the ingress interface.
Invalid Destination Address	The number of OSPF packets discarded because the packet's destination IP address is not the address of the ingress interface and is not the *AllDrRouters* or *AllSpfRouters* multicast addresses.

Field	Description
Wrong Authentication Type	The number of packets discarded because the authentication type specified in the OSPF header does not match the authentication type configured on the ingress interface.
Authentication Failure	The number of OSPF packets dropped because authentication with the neighbor failed. This authentication occurs at the interface level.
No Neighbor at Source Address	The number of OSPF packets dropped because the sender is not an existing neighbor or the sender's IP address does not match the previously recorded IP address for that neighbor.
Invalid OSPF Packet Type	The number of OSPF packets discarded because the packet type field in the OSPF header is not a known type.
Hellos Ignored	The number of received hello packets that were ignored by the interface. These are the hello packets from the new neighbors after the limit was reached for the number of neighbors on the interface or on the switch.
Hellos Sent	The number of hello packets sent on the interface.
Hellos Received	The number of hello packets received on the interface.
DD Packets Sent	The number of database description (DD) packets sent on the interface.
DD Packets Received	The number of database description packets received on the interface.
LS Requests Sent	The number of link state (LS) requests sent on the interface.
LS Requests Received	The number of link state requests received on the interface.
LS Updates Sent	The number of link state updates sent on the interface.
LS Updates Received	The number of link state updates received on the interface.
LS Acknowledgements Sent	The number of link state acknowledgements sent on the interface.
LS Acknowledgements Received	The number of link state acknowledgements received on the interface.

View or clear OSPF neighbor information for an interface

If OSPF is enabled, you can view the OSPF neighbors. You can also clear the OSPF neighbor information.

To view or clear OSPF neighbor information for an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > Neighbor Table**.
   The OSPF Neighbor Table page displays.
3. To refresh the page, click the **Refresh** button.
4. To clear the OSPF neighbor information, resetting the fields in the table to default values, click the **Clear** button.
5. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
Interface	The interface for which data is displayed.
Neighbor IP Address	The IP address of the neighboring router's interface that is attached to the network. The neighbor IP address is learned when hello packets are received from the neighbor. For virtual links, the neighbor IP address is learned during the routing table build process. The neighbor IP address is used in the following situations: • as the destination IP address when protocol packets are sent as unicast packets along this adjacency. • in a router LSAs as the link ID for the attached network if the neighboring router is selected as the designated router.
Neighbor Interface Index	The unit/slot/port combination that identifies the neighbor interface index.
Router ID	A 32-bit integer in dotted-decimal format that represents the neighbor interface.
Area ID	The area ID of the OSPF area associated with the interface.
Options	The integer value that indicates the optional OSPF capabilities supported by the neighbor. The neighbor's optional OSPF capabilities are also listed in its hello packets. The Options information allows received hello packets to be rejected if there is a mismatch in certain OSPF capabilities. (If hello packets are rejected, for example, neighbor relationships cannot be formed.)
Router Priority	The OSPF priority for the interface. The priority of an interface is a priority integer from 0 to 255. A value of 0 indicates that the router is not eligible to become the designated router on this network.
State The state of the OSPD neighbor:	
Down: This is the initial state of a neighbor conversation. It indicates that no recent information was received from the neighbor. On Non-broadcast multiple access (NBMA) networks, hello packets can still be sent to neighbors in the Down state, although at a reduced frequency.	
Attempt: This state is valid only for neighbors attached to NBMA networks. It indicates that no recent information was received from the neighbor, but that a more concerted effort must be made to contact the neighbor. This is done by sending the neighbor hello packets at hello intervals.	
Init: A hello packet was recently detected from the neighbor. However, bidirectional communication is not yet established with the neighbor (for example, the router is not included in the neighbor's hello packet). All neighbors in this state or a state listed below in this table are	

Field	Description
included in the hello packets sent from the associated interface.	
2-Way: Communication between the two routing interfaces is bidirectional. This is assured by the operation of the hello protocol. This is the most advanced state short of beginning adjacency establishment. The backup designated router is selected from the set of neighbors in the 2-Way state or a state listed below in this table.	
Exchange Start: This is the first step in creating an adjacency between the two neighboring routing interfaces. The process in this state determines which router is the master and what the initial DD sequence number is. Neighbor conversations in this state or a state listed below in this table are called adjacencies.	
Exchange: The routing interfaces includes its entire link state database in database description packets that it sends to the neighbor. The router can also send link state request packets to request the neighbor's more recent LSAs. All adjacencies in the Exchange state or a state listed below in this table are used by the flooding procedure. These adjacencies can transmit and receive all types of OSPF routing protocol packets.	
Loading: Link state request packets are sent to the neighbor requesting the most recent LSAs that were discovered but not yet received in the Exchange state.	
Full: The neighboring routers are fully adjacent. These adjacencies are now included in router LSAs and network LSAs.	
Events	The number of times this neighbor relationship changed state, or an error occurred.
Permanence	The status of the entry. Dynamic and Permanence refer to how the neighbor became known.
Hellos Suppressed	This indicates whether hellos are being suppressed to the neighbor.
Retransmission Queue Length	An integer representing the current length of the retransmission queue of the specified neighbor router ID of the specified interface.
Up Time	The neighbor uptime, which indicates the time since the adjacency last reached the Full state.
Dead Time	The time in seconds that the switch waits before determining that the neighbor is unreachable.

View the OSPF link state database

If OSPF is enabled, you can view the OSPF link state database.

To view the OSPF link state database:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > Link State Database**.
   The page displays the Link State Database, External LSDB Table, and AS Opaque LSDB Table sections.
   The term LSDB stands for Link State Database; the term AS stands for autonomous system.
3. To refresh the page, click the **Refresh** button.

The following tables describe the view-only fields on the page.

Table 132.Link State Database

Field	Description
Router ID	A 32-bit number in dotted-decimal format that identifies the switch within the autonomous system (AS). To change the router ID, you must first disable OSPF. After you set a new router ID, you must reenable OSPF for the change to take effect. The default value is 0.0.0.0, although this is not a valid router ID.
Area ID	The ID of the OSPF area to which one of the switch interfaces is connected. An area ID is a 32-bit integer in dotted-decimal format that uniquely identifies the area to which an interface is connected.
LSA Type	The format and function of the link state advertisement: Illegal; Router Links; Network Links; Network Summary; ASBR Summary; AS-external; Group Member; NSSA; TMP2; Link Opaque; Area Opaque; AS Opaque; or Unknown.
LS ID	The link state ID identifies the part of the routing domain that is being described by the advertisement. The value of the LS ID depends on the advertisement's LS type.
Age	The time in seconds since the link state advertisement was first originated.
Sequence	A signed 32-bit integer that is used to detect old and duplicate link state advertisements. The larger the sequence number, the more recent the advertisement.
Checksum	The checksum of the complete contents of the advertisement, except the LS age field. The checksum is used to detect data corruption of an advertisement. This corruption can occur while an advertisement is being flooded, or while it is being held in a router's memory.
Options	The Options field in the link state advertisement header indicates which optional capabilities are associated with the advertisement: • **Q**: Enables support for QoS traffic engineering. • **E**: Specifies the way AS external LSAs are flooded. • **MC**: Specifies the way IP multicast datagrams are forwarded according to the standard specifications. • **O**: Specifies if opaque LSAs are supported. • **V**: Specifies if OSPF extensions for class of service (CoS) and VPN are supported.

Table 133.External LSDB table

Field	Description
Router ID	A 32-bit number in dotted-decimal format that identifies the switch within the autonomous system (AS). To change the router ID, you must first disable OSPF. After you set a new router ID, you must reenable OSPF for the change to take effect. The default value is 0.0.0.0, although this is not a valid router ID.
LSA Type	The format and function of the link state advertisement: ASBR Summary; AS-external; NSSA; or TMP2.
LS ID	The link state ID identifies the part of the routing domain that is being described by the advertisement. The value of the LS ID depends on the advertisement's LS type.
Age	The time in seconds since the link state advertisement was first originated.
Sequence	A signed 32-bit integer that is used to detect old and duplicate link state advertisements. The larger the sequence number, the more recent the advertisement.
Checksum	The checksum of the complete contents of the advertisement, except the LS age field. The checksum is used to detect data corruption of an advertisement. This corruption can occur while an advertisement is being flooded, or while it is being held in a router's memory.

Table 134.AS Opaque LSDB table

Field	Description
Router ID	A 32-bit number in dotted-decimal format that identifies the switch within the autonomous system (AS). To change the router ID, you must first disable OSPF. After you set a new router ID, you must reenable OSPF for the change to take effect. The default value is 0.0.0.0, although this is not a valid router ID.
LSA Type	The format and function of the link state advertisement: Area Opaque; AS Opaque; or Link Opaque.
LS ID	The link state ID identifies the part of the routing domain that is being described by the advertisement. The value of the LS ID depends on the advertisement's LS type.
Age	The time in seconds since the link state advertisement was first originated.

Field	Description
Sequence	A signed 32-bit integer that is used to detect old and duplicate link state advertisements. The larger the sequence number, the more recent the advertisement.
Checksum	The checksum of the complete contents of the advertisement, except the LS age field. The checksum is used to detect data corruption of an advertisement. This corruption can occur while an advertisement is being flooded, or while it is being held in a router's memory.

Configure an OSPF virtual link

If you cannot set up a physical link from the switch to the OSPF backbone area, you can add a virtual link to connect to the backbone area through a common (non-backbone) area. You can configure a virtual link between any area border routers that have physical interfaces to the common area.

To configure an OSPF virtual link:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > Virtual Link Configuration**.
   The OSPF Virtual Link Configuration page displays.
3. In the **Area ID** field, enter an OSPF area ID.
   An area ID is a 32-bit integer in dotted-decimal format that identifies the area to which a switch interface connects.
4. In the **Neighbor Router ID** field, enter the neighbor portion of a virtual link specification.
5. In the **Hello Interval** field, enter the OSPF hello interval for the interface.
   This interval must be the same for all routers and switches in the network. The interval ranges from 1 to 65,535 seconds. The default is 10 seconds.
6. In the **Dead Interval** field, enter the number of seconds that the interface waits for a neighbor's hello packets before determining that the neighbor is down.
   This interval must be the same for all routers and switches in the network. This interval must be a multiple of the hello interval. The interval ranges from 1 to 65,535. The default is 40 seconds. (The default hello interval is 10 seconds.)

7. In the **Iftransit Delay Interval** field, enter the estimated number of seconds required to transmit a link state update packet over the interface.
   The interval range from 1 to 3600 seconds (1 hour). The default is 1 second.
8. In the **Retransmit Interval** field, enter the number of seconds between link-state advertisements for adjacencies, that is, for neighboring devices that form an OSPF adjacency with the interface.
   The interval is also used when retransmitting database descriptions and link-state request packets. The value for the interval ranges from 1 to 3600 seconds (1 hour). The default is 5 seconds.
9. From the **Authentication Type** menu, select if communication over the interface is authenticated and encrypted:
   - **None**: The interface does not use authentication and encryption. This is the default settings.
   - **Simple**: The interface prompts you to enter an authentication key. This key is included (non-encrypted) in the OSPF header of all packets sent on the network. All routers and switches on the network must be configured with the same key.
   - **Encrypt**: The interface prompts you to enter an authentication key and an authentication ID. Encryption uses the MD5 Message-Digest algorithm. All routers and switches on the network must be configured with the same key and ID.
10. In the **Authentication Key** field, enter the OSPF authentication key for the interface.
    If you select **Simple** from the **Authentication Type** menu, you can set a key with maximum of 8 octets. If you select **Encrypt** from the **Authentication Type** menu, you can set a key with maximum of 16 octets.
11. In the **Authentication ID** field, enter the ID that must be used for authentication on the interface.
    If you select **Encrypt** from the **Authentication Type** menu, you must set an ID, which is a number between 0 and 255.
12. Click the **Add** button.
    Your settings are saved.
13. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
Neighbor State The state of the OSPF neighbor interface:	
Down: This is the initial interface state. The interface is not usable for OSPF, and the interface settings are at their initial values. All interface timers are disabled, and no adjacencies are associated with the interface.	
Loopback: The interface is looped back either in hardware or software. The interface is unavailable for regular data traffic. You can get information about the interface by sending ICMP pings to the interface or through a bit error test because IP packets can still be addressed to an interface in a loopback state. An interface in a loopback state is advertised in LSAs as single host route for which the destination is the IP interface address.	
Waiting: The neighboring device is monitoring received hello packets to detect the identity of the backup designated router for the network. The neighboring device cannot elect a backup designated router or a designated router until it transitions out of the waiting state. This prevents unnecessary changes of the backup designated router.	
Designated Router: The neighboring device is the designated router on the attached network. Adjacencies are established to all other routers and switches that in the network. The neighboring device must also originate a network LSA for the network node. The network LSA contains links to all switches and routers (including the designated router) in the network.	
Backup Designated Router: The neighboring device is the backup designated router on the attached network. The neighboring device is promoted to designated router if the current designated router fails. The neighboring device establishes adjacencies to all other routers attached to the network. The backup designated router performs slightly different functions during the LSA flooding, as compared to the designated router.	
Other Designated Router: The interface is connected to a broadcast on which other routers are the designated router and backup designated router. The neighboring device attempts to form adjacencies to both the designated router and the backup designated router.	
State The state of the OSPF switch interface:	
Down: This is the initial interface state. The interface is not usable for OSPF, and the interface settings are at their initial values. All interface timers are disabled, and no adjacencies are associated with the interface.	
Loopback: The interface is looped back either in hardware or software. The interface is unavailable for regular data traffic. You can get information about the interface by sending ICMP pings to the interface or through a bit error test because IP packets can still be addressed to an interface in a loopback state. An interface in a loopback state is advertised in LSAs as single host	

Field	Description
route for which the destination is the IP interface address.	
Waiting: The switch is monitoring received hello packets to detect the identity of the backup designated router for the network. The switch cannot elect a backup designated router or a designated router until it transitions out of the waiting state. This prevents unnecessary changes of the backup designated router.	
Designated Router: The switch is the designated router on the attached network. Adjacencies are established to all other routers and switches that in the network. The switch must also originate a network LSA for the network node. The network LSA contains links to all switches and routers (including the designated router) in the network.	
Backup Designated Router: The switch is the backup designated router on the attached network. The switch is promoted to designated router if the current designated router fails. The switch establishes adjacencies to all other routers attached to the network. The backup designated router performs slightly different functions during the LSA flooding, as compared to the designated router.	
Other Designated Router: The interface is connected to a broadcast on which other routers are the designated router and backup designated router. The switch attempts to form adjacencies to both the designated router and the backup designated router.	
Metric	The metric value that is used by the virtual link to compute the shortest paths.

Configure the OSPF route redistribution

Route redistribution (RR) lets the switch share routing information between locally connected routers (or routing interfaces) and between different routing protocols, such as OSPF and RIP, as well as information about static routes.

To configure the OSPF route redistribution:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > Route Redistribution**.
   The OSPF Route Redistribution page displays.

3. Next to the **Source** field, select the check box for one or more types of source routes to which the configuration applies:
   - **Connected**: Routes between locally connected routers (or routing interfaces).
   - **Static**: Routes that were manually configured.
   - **RIP**: Routes that were detected by Routing Information Protocol (RIP).
4. From the **Redistribute** menu, select if redistribution is enabled:
   - **Enable**: The type of route can be redistributed.
   - **Disable**: The type of route cannot be redistributed.
5. In the **Metric Value** field, set the metric value for the type of route.
   If the type of route is preconfigured, the metric value displays, but you can change it. The range is from 0 to 16777214.
6. From the **Metric Type** menu, select the type of metric that must used for the type of route:
   - **External Type 1**: External type 1 metrics that are comparable to the OSPF metric.
   - **External Type 2**: External type 2 metrics that are assumed to be larger than the cost of the OSPF metric.
7. In the **Tag** field, set the tag for the type of route.
   If the type of route is preconfigured, the tag displays, but you can change it. The range is from 0 to 4294967295.
8. From the **Subnets** menu, select if subnetted routes can be redistributed:
   - **Enable**: Subnetted routes can be redistributed.
   - **Disable**: Subnetted routes cannot be redistributed.
9. In the **Distribute List** field, specify an existing access control list (ACL) that filters the routes to be redistributed by the destination protocol.
   Only permitted routes are redistributed. If you enter a nonexistent ACL, all routes are permitted. The range for ACL IDs is from 1 to 199.
   For route filtering, only the following ACL fields are used:
   - **Source IP address and netmask**: The source IP address is compared to the destination IP address of the route. The source IP netmask is treated as a wildcard mask, indicating which bits in the source IP address must match the destination address of the route. (A 1 in the mask indicates a *do not care* in the corresponding address bit.)
   - **Destination IP address and netmask**: The destination IP address is compared to the source IP address of the route. The destination IP netmask is treated as a wildcard mask, indicating which bits in the destination IP address must match the source address of the route. (A 1 in the mask indicates a *do not care* in the corresponding address bit.)
   - **Action**: Permit or Deny.

   All other fields in the ACL (source and destination port, precedence, ToS, and so on) are ignored.

10. Click the **Apply** button.
    Your settings are saved.
11. To save the settings to the running configuration, click the **Save** icon.

Configure OSPF nonstop forwarding

OSPF nonstop forwarding (NSF) ensures that OSPF routing continues by temporarily letting other devices in the network perform OSPF routing for the switch. That means that even if the switch restarts, OSPF routing can continue. (We refer to this type of restart process as a graceful restart.) You can configure if OSPF NSF is enabled on the switch and if the switch can support OSPF NSF events on neighbors, that is, help neighbours to gracefully restart.

To configure OSPF NSF:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPF > Advanced > NSF OSPF Summary**.
   The NSF OSPF Summary page displays.
3. From the **Support Mode** menu, select how the switch performs a graceful restart:
   - **Always**: OSPF performs a graceful restart for all planned and unplanned warm restart events.
   - **Disabled**: OSPF graceful restarts are disabled. This is the default selection.
   - **Planned**: OSPF performs a graceful restart only if a restart is planned (for example, because a CLI command is executed).
4. In the **Restart Interval** field, enter the time that the switch waits after a graceful restart before continuing OSPF routing.
   The range is from 0 to 1800 seconds. The default is 120 seconds.
5. From the **Helper Support Mode** menu, select if the switch can support a neighbor that performs a warm restart:
   - **Always**: OSPF helps a restarting neighbor for all planned and unplanned warm restart events. This is the default selection.
   - **Disabled**: OSPF does not help a restarting neighbor.
   - **Planned**: OSPF helps a restarting neighbor only if the restart is planned.
6. From the **Helper Strict LSA Checking** menu, select if the switch exits helper support mode when the topology changes:
   - **Enable**: The switch exits helper support mode when the topology changes. This is the default selection.
   - **Disable**: The switch does not exit helper support mode when the topology changes.

7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Field	Description
Restart Status	The restart status of the OSPF helper support mode: • Not Restarting • Planned Restart • Unplanned Restart
Restart Age (secs)	The time in seconds since the last restart occurred.
Restart Exit Reason	Displays how the OSPF routing process on the switch last started: • **Not Attempted**: A graceful restart was not attempted. • **Progress**: A graceful restart is in progress. • **Completed**: The previous graceful restart completed successfully. • **Timed Out**: The previous graceful restart timed out. • **Topology Changed**: The previous graceful restart terminated prematurely because of a topology change.

Open Shortest Path First version 3

Open Shortest Path First version 3 (OSPFv3) is the OSPF routing protocol for IPv6. It is similar to OSPFv2 (generally just called OSPF) in its concept of a link state database, intra- and inter-area, and AS external routes and virtual links.
You can use OSPF and OSPFv3 simultaneously: OSPF works with IPv4, and OSPFv3 works with IPv6.
OSPFv3 differs from OSPF in a number of respects:

- Peering is done through link-local addresses
- The protocol is link based rather than network based
- Address changes are communicated through leaf Link State Advertisements (LSAs), which eventually allow its use for both IPv4 and IPv6
- Point-to-point links are supported to enable operation over tunnels

Enable OSPFv3

Before OSPFv3 can become operational, you must enable the administrative mode for IPv6 routing on the switch (see Manage the global IPv6 routing settings on page 363).

To enable OSPFv3:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Basic > OSPFv3 Configuration**.
   The OSPFv3 Configuration page displays.
3. Select an Admin Mode radio button:
   - **Enable**: OSPFv3 is enabled for the switch. This is the default setting.
   - **Disable**: OSPFv3 is disabled for the switch.

   The Router ID field displays the 32-bit number in dotted-decimal format. This number uniquely identifies the switch within the autonomous system (AS).
   To change the router ID, you must first disable OSPFv3. After you set a new router ID, you must reenable OSPFv3 for the change to take effect. The default value is 0.0.0.0, although this is not a valid router ID.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the OSPFv3 default route advertisement

Whether or not a default route exists in the routing table, you can set up a default route for OSPFv3 advertisements, including a metric value and a type of metric.

To configure the OSPFv3 default route advertisement:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Advanced > OSPFv3 Configuration**.
   The OSPFv3 Configuration page displays. The page shows different sections.
3. Go to the Default Route Advertise Configuration section.
4. Select a Default Information Originate radio button:
   - **Enable**: OSPFv3 originates a default route for external link-state advertisements (LSAs). This route is in the format IPv6 address/prefix.
   - **Disable**: OSPFv3 does not originate a default route for external LSAs. This is the default settings.

5. Select an Always radio button:
   - **True**: OSPFv3 originates a default route regardless of whether a default route exists.
   - **False**: OSPFv3 originates a default route only if a default route is already in the switch's routing table.
6. In the **Metric** field, set the metric value of the default route.
   The value ranges from 0 to 16777214. The default is 0.
7. Select a Metric Type field radio button:
   - **External Type 1**: This metric is computed by adding the internal OSPFv3 cost to the external redistributed cost.
   - **External Type 2**: This metric is computed by using only the external redistributed cost. This is the default.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Configure the global OSPFv3 settings

You can configure the OSPFv3 settings that apply globally to the OSPFv3 processes on the switch.

To configure the global OSPFv3 settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Advanced > OSPFv3 Configuration**.
   The OSPFv3 Configuration page displays. The page shows different sections.
3. Go to the OSPFv3 Configuration section.
4. In the **Router ID** field, enter the 32-bit number in dotted-decimal format that uniquely identifies the router within the autonomous system (AS).
   To change the router ID, you must first disable OSPFv3. After you set a new router ID, you must reenable OSPFv3 for the change to take effect. The default value is 0.0.0.0, although this is not a valid router ID.
5. Select an Admin Mode radio button:
   - **Enable**: OSPFv3 is enabled for the switch. This is the default setting.
   - **Disable**: OSPFv3 is disabled for the switch.

   Before OSPFv3 can become operational, you must enable the administrative mode for IPv6 routing on the switch (see Configure IPv6 routing interfaces on page 365).

6. In the **Exit Overflow Interval** field, set how many seconds OSPFv3 must wait before leaving an overflow state.
   When the number of nondefault external LSAs exceeds a limit, the switch enters an overflow state in which OSPFv3 cannot originate nondefault external LSAs. (For more information, see RFC 1765.)
   If the number of seconds is 0 (the default setting), OSPFv3 does not leave the overflow state until you disable and reenable OSPFv3. The maximum value can be 2,147,483,647 seconds.
7. In the **External LSDB Limit** field, set the limit of the external link-state database (LSDB) for OSPFv3.
   The router's link state database cannot hold more nondefault AS-external-LSAs than the limit that is set for the external LSDB.
   The external LSDB's limit must be set identically on all routers that are attached to the OSPFv3 backbone or any regular OSPFv3 area. The limit can be from –1 to 2147483647. The default setting is –1, which indicates that there is no limit.
8. In the **Default Metric** field, set the default for the metric of redistributed routes.
   The range is 1 from 16777214. The default is 0.
9. In the **Maximum Paths** field, set the number of paths that OSPFv3 can report for a destination.
   The range is from 1 to 16. The default is 16.
10. In the **AutoCost Reference Bandwidth** field, set the bandwidth that controls how OSPFv3 calculates link cost.
    Set the reference bandwidth in megabits per second. Unless a link cost is configured, the link cost is computed by dividing the reference bandwidth by the interface bandwidth. The range is from 1 to 4294967. The default is 100.
11. From the **Default Passive Setting** menu, configure the global passive mode setting for all OSPFv3 interfaces:
    - **Enable**: The global passive mode setting is enabled. OSPFv3 does not form adjacencies on passive interfaces, but does advertise attached networks as stub networks.
    - **Disable**: The global passive mode setting is disabled. The default is Disable.
12. From the **Helper Support Mode** menu, select if the switch can support a neighbor that performs a warm restart in an OSPFv3 nonstop forwarding (NSF) configuration:
    - **Always**: OSPFv3 helps a restarting neighbor for all planned and unplanned warm restart events. This is the default selection.
    - **Disabled**: OSPFv3 does not help a restarting neighbor.
    - **Planned**: OSPFv3 helps a restarting neighbor only if the restart is planned.

13. From the **Helper Strict LSA Checking** menu, select if the switch exits helper support mode (for OSPFv3 NSF) when the topology changes:
    - **Enable**: The switch exits helper support mode when the topology changes. This is the default selection.
    - **Disable**: The switch does not exit helper support mode when the topology changes.
14. Click the **Apply** button.
    Your settings are saved.
15. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
ASBR Mode	Indicates if the routing interface functions as an autonomous system boundary router (ASBR). An ASBR can be configured to redistribute routes from another protocol or to originate external LSAs that advertise the default route.
ABR Status	Indicates if the routing interface functions as an area border router (ABR). An ABR has active non-virtual interfaces in two or more OSPFv3 areas.
External LSA Count	The number of external (LS type 5) link state advertisements (LSAs) in the link state database.
External LSA Checksum	The sum of the LS checksums of the external LSAs that are contained in the link state database. You can use this sum to determine if there was a change in the link state database of a switch or router, and to compare the link state databases of two devices. This value is in hexadecimal.
New LSAs Originated	The number of LSAs originated by the switch. In any OSPFv3 area, a router (or in this situation, a switch) originates several LSAs. Each router originates a router-LSA. If the router is also the designated router for any of the area's networks, it originates network LSAs for those networks.
LSAs Received	The number of received LSAs for new instantiations. This number does not include new instantiations of self-originated LSAs.

Add or delete an OSPFv3 common area ID

An OSPFv3 common area consists of a group of devices that share the same area ID.

To add or delete an OSPFv3 common area ID:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Advanced > Common Area Configuration**.
   The OSPFv3 Common Area Configuration page displays.
3. In the **Area ID** field, enter an OSPFv3 area ID.
   An area ID is a 32-bit integer in dotted-decimal format that identifies the area to which a switch interface connects.
4. Take one of the following actions:
   - **Add**: Click the **Add** button. The area ID added.
   - **Delete**: Click the **Delete** button. The area ID deleted.
5. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
External Routing	A definition of the switch's capabilities for the area, including whether or not AS-external-LSAs are flooded into or throughout the area. If the area is a stub area, one of the following external routing capabilities can display: • **Import External LSAs**: Imports and propagates external LSAs. • **Import No LSAs**: Does not import and propagate external LSAs. If the area is *not* a stub area, the only option is Import No LSAs.
SPF Runs	The number of times that the intra-area route table was calculated using this area's link state database. (This is typically calculated using Dijkstra's algorithm.)
Area Border Router Count	The total number of area border routers that re reachable within this area. This is initially zero, and is calculated in each SPF pass.
Area LSA Count	The total number of link state advertisements in this area's link state database, excluding AS external LSAs.
Area LSA Checksum	The 32-bit unsigned sum of the link state advertisement checksum contained in this area's link state database. This sum excludes external (LSA type 5) link state advertisements. You can use the sum to determine if a change occurred in the link state database of the switch, and to compare the link state database of two devices.
Import Summary LSAs	The summary LSAs that are imported into this area.

Add an OSPFv3 stub area

A stub area is a controlled area that does not allow external routes from outside the OSPFv3 network.

To add an OSPFv3 stub area:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Advanced > Stub Area Configuration**.
   The OSPFv3 Stub Area Configuration page displays.
3. In the **Area ID** field, enter an OSPFv3 area ID.
   An area ID is a 32-bit integer in dotted-decimal format that identifies the area to which a switch interface connects.
4. From the **Import Summary LSAs** menu, select if summary LSAs are imported:
   - **Enable**: Summary LSAs are imported into stub areas.
   - **Disable**: Summary LSAs are not imported into stub areas.
5. In the **Default Cost** field, enter the metric value that must be applied for the default route advertised to the stub area.
   The range is from 1 to 16,777,215.
6. Click the **Add** button.
   The area is added as a stub area.
7. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
SPF Runs	The number of times that the intra-area route table was calculated using this area's link state database. (This is typically calculated using Dijkstra's algorithm.)
Area Border Router Count	The total number of area border routers that are reachable within this area. This is initially zero, and is calculated in each SPF pass.
Area LSA Count	The total number of link state advertisements in this area's link state database, excluding AS external LSAs.
Area LSA Checksum	The 32-bit unsigned sum of the link state advertisement checksum contained in this area's link state database. This sum excludes external (LSA type 5) link state advertisements. You can use the sum to determine if a change occurred in the link state database of the switch, and to compare the link state database of two devices.
Type of Service	The normal ToS that is associated with the stub metric.

Add an OSPFv3 NSSA area

The OSPFv3 not-so-stubby area (NSSA) can receive external routes (Type 7 LSA) and forward these routes to other areas. You can select if summary LSAs (Type 3 LSAs) can be imported into stub areas.

To add an OSPFv3 NSSA area:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Advanced > NSSA Area Configuration**.
   The OSPFv3 NSSA Area Configuration page displays.
3. In the **Area ID** field, enter an OSPFv3 area ID.
   An area ID is a 32-bit integer in dotted-decimal format that identifies the area to which a switch interface connects.
4. From the **Import Summary LSAs** menu, select if summary LSAs are imported:
   - **Enable**: Summary LSAs are imported into stub areas.
   - **Disable**: Summary LSAs are not imported into stub areas.

5. The options in the Default Information Originate section of the table let you configure the switch to advertise a default route into the NSSA when the import of summary LSAs is disabled:
   a. From the **Admin Mode** menu, select **Enable** to let the switch advertise a default route or **Disable** to prevent the switch from doing so.
   b. In the **Metric Value** field, set the metric value for the default route. The range is from 1 to 16777214.
   c. From the **Metric Type** menu, select the type of metric that must used in the default route:
      - **Comparable Cost**: External type 1 metrics that are comparable to the OSPFv3 metric
      - **Non-comparable Cost**: External type 2 metrics that are assumed to be larger than the cost of the OSPFv3 metric.
6. From the **Translator Role** menu, select the role of the NSSA:
   - **Always**: The switch assumes the role of the translator when it becomes a border router.
   - **Candidate**: The switch participates in the translator election process when it attains border router status.
7. In the **Translator Stability Interval** field, enter the period during which the switch extends its translator role after another device becomes the translator.
   The range is from 0 to 3600 minutes.
8. From the **Redistribute Mode** menu, select if learned external routes are redistributed:
   - **Enable**: The NSSA ABR distributes learned external routes to the NSSA.
   - **Disable**: The NSSA ABR does not distribute learned external routes to the NSSA.
9. Click the **Add** button.
   The area is added as an NSSA area.
10. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
SPF Runs	The number of times that the intra-area route table was calculated using this area's link state database. (This is typically calculated using Dijkstra's algorithm.)
Area Border Router Count	The total number of area border routers that re reachable within this area. This is initially zero, and is calculated in each SPF pass.
Area LSA Count	The total number of link state advertisements in this area's link state database, excluding AS external LSAs.
Area LSA Checksum	The 32-bit unsigned sum of the link state advertisement checksum contained in this area's link state database. This sum excludes external (Type 5 LSAs) link state advertisements. You can use the sum to determine if a change occurred in the link state database of the switch, and to compare the link state database of two devices.
Translator State	This field displays if and how the NSSA border router translates Type 7 LSAs into Type 5 LSAs: • **Enabled**: The NSSA border router's translator role is set to always. • **Elected**: The candidate NSSA border router is translating Type 7 LSAs into Type 5 LSAs. • **Disabled**: The candidate NSSA border router is not translating Type 7 LSAs into Type 5 LSAs.

Add an OSPFv3 area range

You add an OSPFv3 area range by setting an area ID with an address range (an IPv6 prefix). You then also set the LSAs that are associated with the area and whether the address range is advertised outside the area.
You can associate an OSPFv3 area range with a common area (see Add or delete an OSPFv3 common area ID on page 440), stub area (see Add an OSPFv3 stub area on page 442), or NSSA area (see Add an OSPFv3 NSSA area on page 443).

To add an OSPFv3 area range:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Advanced > Area Range Configuration**.
   The OSPFv3 Area Range Configuration page displays.

3. In the **Area ID** field, enter an OSPFv3 area ID.
   An area ID is a 32-bit integer in dotted-decimal format that identifies the area to which a switch interface connects.
4. In the **IPv6 Prefix** field, enter the IPv6 address and prefix for the selected area.
5. From the **LSDB Type** menu, select the type of link state advertisements (LSAs) that must be associated with the area ID and address range.
   - **Network Summary**: The switch generates Network Summary LSAs (Type 3 LSAs) for the area. This is the default setting.
   - **NSSA External**: The switch generates NSSA External LSAs (Type 7 LSAs) for the area. These types of LSAs provide more information.
6. From the **Advertise** menu, select if the address range is advertised outside the area:
   - **Enable**: The address range is advertised outside the area through Network Summary LSAs. This is the default setting.
   - **Disable**: The address range is not advertised outside the area.
7. Click the **Add** button.
   The OSPFv3 area ID and range is added.
8. To save the settings to the running configuration, click the **Save** icon.

Configure an OSPFv3 interface

You can configure one or more OSPFv3 interfaces, which are interfaces that are connected to OSPFv3 areas over point-to-point or broadcast connections.

To configure an OSPFv3 interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Advanced > Interface Configuration**.
   The OSPFv3 Interface Configuration page displays.

3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces or VLANs by taking one of the following actions:
   - To configure a single interface or VLAN, select the check box associated with the interface or VLAN, or type the interface or VLAN number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces or VLANs with the same settings, select the check box associated with each interface or VLAN.
5. In the **Area ID** field, enter an OSPFv3 area ID.
   An area ID is a 32-bit integer in dotted-decimal format that identifies the area to which a switch interface connects.
6. Select an Admin Mode radio button:
   - **Enable**: OSPFv3 is enabled for the interface.
   - **Disable**: OSPFv3 is disabled for the interface. This is the default setting.

   OSPFv3 can become operational on the interface only if the interface is configured with an IP address and subnet mask (see Configure IPv6 routing interfaces on page 365).
   You can configure the OSPFv3 settings without enabling OSPFv3 on the interface, but the change does not take effect until you enable OSPFv3 on the interface. The following fields display only if you enable OSPFv3 on the interface:
   State, Designated Router, Backup Designated Router, and Number of Link Events.
7. In the **Router Priority** field, enter the OSPFv3 priority for the interface.
   This is a number from 0 to 255. The default is 1, which is the highest router priority. A value of 0 indicates that the router is not eligible to become the designated router on this network.
8. In the **Retransmit Interval** field, enter the number of seconds between link-state advertisements for adjacencies, that is, for neighboring devices that form an OSPFv3 adjacency with the interface.
   The interval is also used when retransmitting database descriptions and link-state request packets. The value for the interval ranges from 1 to 3600 seconds (1 hour). The default is 5 seconds.

9. In the **Hello Interval** field, enter the OSPFv3 hello interval for the interface.
   This interval must be the same for all routers and switches in the network. The interval ranges from 1 to 65,535 seconds. The default is 10 seconds.
10. In the **Dead Interval** field, enter the number of seconds that the interface waits for a neighbor's hello packets before determining that the neighbor is down.
    This interval must be the same for all routers and switches in the network. This interval must be a multiple of the hello interval. The interval ranges from 1 to 65,535. The default is 40 seconds. (The default hello interval is 10 seconds.)
11. In the **Iftransit Delay Interval** field, enter the estimated number of seconds required to transmit a link state update packet over the interface.
    The interval range from 1 to 3600 seconds (1 hour). The default is 1 second.
12. From the In the **MTU Ignore** menu, select if OSPFv3 MTU mismatch detection is disabled:
    - **Enable**: OSPFv3 MTU mismatch detection is *disabled* on incoming database description packets. The default value is Disable (MTU mismatch detection is enabled).
    - **Disable**: OSPFv3 MTU mismatch detection is *enabled* on incoming database description packets. This is the default setting.
13. From the **Passive Mode** menu, select if the interface does not form an adjacency:
    - **Enable**: The interface is passive to prevent OSPFv3 from forming an adjacency. OSPFv3 advertises networks attached to passive interfaces as stub networks.
    - **Disable**: The interface is not passive and OSPFv3 can form an adjacency. This is the default setting.
14. From the **Network Type** menu, select the type of OSPFv3 network for the interface:
    - **Broadcast**: The interface can connect to multiple OSPFv3 routers. This is the default setting.
    - **Point to point**: The interface can connect to one other OSPFv3 router.
15. In the **Metric Cost** field, enter the value that OSPFv3 uses to compute the shortest paths.
    The range is from 1 to 65,535. The default is 1.
16. Click the **Apply** button.
    Your settings are saved.
17. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
IPv6 Address	The IP address of the interface
LSA Ack Interval	The number of seconds that the interface must wait before sending a delayed acknowledgement.
State The state of the OSPFv3 interface:	
Down: This is the initial interface state. The interface is not usable for OSPFv3, and the interface settings are at their initial values. All interface timers are disabled, and no adjacencies are associated with the interface.	
Loopback: The interface is looped back either in hardware or software. The interface is unavailable for regular data traffic. You can get information about the interface by sending ICMP pings to the interface or through a bit error test because IP packets can still be addressed to an interface in a loopback state. An interface in a loopback state is advertised in LSAs as single host route for which the destination is the IP interface address.	
Waiting: The switch is monitoring received hello packets to detect the identity of the backup designated router for the network. The switch cannot elect a backup designated router or a designated router until it transitions out of the waiting state. This prevents unnecessary changes of the backup designated router.	
Designated Router: The switch is the designated router on the attached network. Adjacencies are established to all other routers and switches that in the network. The switch must also originate a network LSA for the network node. The network LSA contains links to all switches and routers (including the designated router) in the network.	
Backup Designated Router: The switch is the backup designated router on the attached network. The switch is promoted to designated router if the current designated router fails. The switch establishes adjacencies to all other routers attached to the network. The backup designated router performs slightly different functions during the LSA flooding, as compared to the designated router.	
Other Designated Router: The interface is connected to a broadcast on which other routers are the designated router and backup designated router. The switch attempts to form adjacencies to both the designated router and the backup designated router.	
Designated Router	The identity of the designated router for this network, as detected by the advertising router. The designated router is identified by its router ID. The value 0.0.0.0 means that no designated router is detected.
Backup Designated Router	The identity of the backup designated router for this network, as detected by the advertising router. The backup designated router is identified by its

Field	Description
	router ID. The value 0.0.0.0 means that no backup designated router is detected.
Number of Link Events	The number of times the OSPFv3 interface changed its state.

View or clear OSPFv3 statistics for an interface

If OSPFv3 is enabled, you can view or clear OSPFv3 statistics for an interface.

To view or clear OSPFv3 statistics for an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Advanced > Interface Statistics**.
   The OSPFv3 Interface Statistics page displays.
3. From the **Interface** menu, select the interface.
4. To refresh the page, click the **Refresh** button.
5. To clear all the counters, resetting all OSPFv3 interface statistics to default values, click the **Clear** button.
6. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
OSPFv3 Area ID	The OSPFv3 area to which the interface belongs. An OSPFv3 area ID is a 32-bit integer in dotted-decimal format that uniquely identifies the area to which the interface connects.
Area Border Router Count	The number of area border routers (ABR) reachable within this area. This is initially zero, and is calculated in each Shortest Path First (SPF) pass.
AS Border Router Count	The total number of autonomous system (AS) border routers reachable within this area. This is initially zero, and is calculated in each SPF pass.
Area LSA Count	The number of link state advertisements in this area's link state database, excluding AS external LSAs.
IPv6 Address	The IPv6 address of the interface.
Interface Events	The number of times the OSPFv3 interface changed its state, or an error occurred.
Virtual Events	The number of state changes or errors that occurred on this virtual link.
Neighbor Events	The number of times this neighbor relationship changed state, or an error occurred.
Sent Packets	The number of OSPFv3 packets transmitted on the interface.
Received Packets	The number of valid OSPFv3 packets received on the interface.
Discards	The number of received OSPFv3 packets discarded because of an error in the packet or an error in processing the packet.
Bad Version	The number of received OSPFv3 packets for which the version field in the OSPFv3 header does not match the version of the OSPFv3 process handling the packet.
Virtual Link Not Found	The number of received OSPFv3 packets discarded because the ingress interface is in a non-backbone area and the OSPFv3 header identified the packet as belonging to the backbone, but OSPFv3 does not have a virtual link to the packet's sender.
Area Mismatch	The number of OSPFv3 packets discarded because the area ID in the OSPFv3 header is not the area ID configured on the ingress interface.
Invalid Destination Address	The number of OSPFv3 packets discarded because the packet's destination IP address is not the address of the ingress interface and is not the *AllDrRouters* or *AllSpfRouters* multicast addresses.
No Neighbor at	The number of OSPFv3 packets dropped because the sender is not an

Field	Description
Source Address	existing neighbor or the sender's IP address does not match the previously recorded IP address for that neighbor.
Invalid OSPF Packet Type	The number of OSPFv3 packets discarded because the packet type field in the OSPFv3 header is not a known type.
Hellos Ignored	The number of received hello packets that were ignored by the interface. These are the hello packets from the new neighbors after the limit was reached for the number of neighbors on the interface or on the switch.
Hellos Sent	The number of hello packets sent on the interface.
Hellos Received	The number of hello packets received on the interface.
DD Packets Sent	The number of database description (DD) packets sent on the interface.
DD Packets Received	The number of database description packets received on the interface.
LS Requests Sent	The number of link state (LS) requests sent on the interface.
LS Requests Received	The number of link state requests received on the interface.
LS Updates Sent	The number of link state updates sent on the interface.
LS Updates Received	The number of link state updates received on the interface.
LS Acknowledgements Sent	The number of link state acknowledgements sent on the interface.
LS Acknowledgements Received	The number of link state acknowledgements received on the interface.

View or clear OSPFv3 neighbor information for an interface

If OSPFv3 is enabled, you can view the OSPFv3 neighbors. You can also clear the OSPFv3 neighbor information.

To view or clear OSPFv3 neighbor information for an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Routing > OSPFv3 > Advanced > Neighbor Table**.
   The OSPFv3 Neighbor Table page displays.
3. To refresh the page, click the **Refresh** button.
4. To clear the OSPFv3 neighbor information, resetting the fields in the table to default values, click the **Clear** button.
5. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
Interface	The interface for which data is displayed.
Interface Identifier	The interface ID that the neighboring router's interface advertises in its hello packets.
Router ID	A 32-bit integer in dotted-decimal format that represents the neighbor interface.
Area ID	The area ID of the OSPFv3 area associated with the interface.
Options	The integer value that indicates the optional OSPFv3 capabilities supported by the neighbor. The neighbor's optional OSPFv3 capabilities are also listed in its hello packets. The Options information allows received hello packets to be rejected if there is a mismatch in certain OSPFv3 capabilities. (If hello packets are rejected, for example, neighbor relationships cannot be formed.)
Router Priority	The OSPFv3 priority for the interface. The priority of an interface is a priority integer from 0 to 255. A value of 0 indicates that the router is not eligible to become the designated router on this network.
State The state of the OSPD neighbor:	
Down: This is the initial state of a neighbor conversation. It indicates that no recent information was received from the neighbor. On Non-broadcast multiple access (NBMA) networks, hello packets can still be sent to neighbors in the Down state, although at a reduced frequency.	
Attempt: This state is valid only for neighbors attached to NBMA networks. It indicates that no recent information was received from the neighbor, but that a more concerted effort must be made to contact the neighbor. This is done by sending the neighbor hello packets at hello intervals.	
Init: A hello packet was recently detected from the neighbor. However, bidirectional communication is not yet established with the neighbor (for example, the router is not included in the neighbor's hello packet). All neighbors in this state or a state listed below in this table are included in the hello packets sent from the associated interface.	
2-Way: Communication between the two routing interfaces is bidirectional. This is assured by the operation of the hello protocol. This is the most advanced state short of beginning adjacency establishment. The backup designated router is selected from the set of neighbors in the 2-Way state or a state listed below in this table.	
Exchange Start: This is the first step in creating an adjacency between the two neighboring routing interfaces. The process in this state determines which router is the master and what the initial DD sequence number is. Neighbor conversations in this state or a state listed below in this table are	

Field	Description
called adjacencies.	
Exchange: The routing interfaces includes its entire link state database in database description packets that it sends to the neighbor. The router can also send link state request packets to request the neighbor's more recent LSAs. All adjacencies in the Exchange state or a state listed below in this table are used by the flooding procedure. These adjacencies can transmit and receive all types of OSPFv3 routing protocol packets.	
Loading: Link state request packets are sent to the neighbor requesting the most recent LSAs that were discovered but not yet received in the Exchange state.	
Full: The neighboring routers are fully adjacent. These adjacencies are now included in router LSAs and network LSAs.	
Dead Time	The time in seconds that the switch waits before determining that the neighbor is unreachable.
Events	The number of times this neighbor relationship changed state, or an error occurred.
Retransmission Queue Length	An integer representing the current length of the retransmission queue of the specified neighbor router ID of the specified interface.

View the OSPFv3 link state database

If OSPFv3 is enabled, you can view the OSPFv3 link state database.

To view the OSPFv3 link state database:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Advanced > Link State Database**.
   The page displays the OSPFv3 Link State Database and OSPFv3 External LSDB Table sections.
   The term LSDB stands for Link State Database.
3. To refresh the page, click the **Refresh** button.

The following tables describe the view-only fields on the page.

Table 135.OSPFv3 Link State Database

Field	Description
Router ID	A 32-bit number in dotted-decimal format that identifies the switch within the autonomous system (AS). To change the router ID, you must first disable OSPFv3. After you set a new router ID, you must reenable OSPFv3 for the change to take effect. The default value is 0.0.0.0, although this is not a valid router ID.
Area ID	The ID of the OSPFv3 area to which one of the switch interfaces is connected. An area ID is a 32-bit integer in dotted-decimal format that uniquely identifies the area to which an interface is connected.
LSA Type	The format and function of the link state advertisement: • **Router LSA**: A router can originate one or more Router LSAs for an area. A Router LSA originates in an area and describes the collected states of all the router's interfaces connected to the area. • **Network LSA**: A Network LSA is originated by the designated router for each link to which two or more routers are attached. The LSA lists all the routers attached to the link. • **Inter-Area Router LSA**: This type LSA is originated by an Area Border Router (ABR) and describes a prefix that is external to the area, yet internal to the autonomous system (AS). • **AS-External LSA**: This type of LSA is originated by an Autonomous System Border Router (ASBR) and describes a path to a prefix external to the autonomous system (AS). • **Link LSA**: A router originates a separate Link LSA for each attached link. It provides the router's link local address to routers that are attached to the link as well as a list of IPv6 prefixes that can be associated with the link. • **Intra-Area-Prefix LSA**: A link's designated router originates one or more Intra-Area-Prefix LSAs to advertise the link's prefixes throughout the area. A router can originate multiple Intra-Area-Prefix LSAs for an area to advertise its own prefixes and those of its attached stub links.
LS ID	The link state ID identifies the part of the routing domain that is being described by the advertisement. The value of the LS ID depends on the advertisement's LS type.
Age	The time in seconds since the link state advertisement was first originated.
Sequence	A signed 32-bit integer that is used to detect old and duplicate link state advertisements. The larger the sequence number, the more recent the

Field	Description
	advertisement.
Checksum	The checksum of the complete contents of the advertisement, except the LS age field. The checksum is used to detect data corruption of an advertisement. This corruption can occur while an advertisement is being flooded, or while it is being held in a router's memory.
Options	The Options field in the link state advertisement header indicates which optional capabilities are associated with the advertisement: • **Q**: Enables support for QoS traffic engineering. • **E**: Specifies the way AS external LSAs are flooded. • **MC**: Specifies the way IP multicast datagrams are forwarded according to the standard specifications. • **O**: Specifies if opaque LSAs are supported. • **V**: Specifies if OSPF extensions for class of service (CoS) and VPN are supported.
Router Options	The Router Options field in the link state advertisement header indicates which type of OSPFv3 router is associated with the advertisement: • **Bit V**: The router is an endpoint of one or more fully adjacent virtual links for which the described area is the transit area. (V stands for virtual link endpoint.) • **Bit E**: The router is an AS boundary router. (E stands for external.) • **Bit B**: The router is an area border router. (B stands for border.) • **Bit Nt**: The router is an NSSA border router that is unconditionally translating NSSA-LSAs into AS-external-LSAs. (Nt stands for NSSA translation.). This means that the router has its NSSA translator role area configuration set to Always.

Table 136.OSPFv3 External LSA Database

Field	Description
Router ID	A 32-bit number in dotted-decimal format that identifies the switch within the autonomous system (AS). To change the router ID, you must first disable OSPFv3. After you set a new router ID, you must reenable OSPFv3 for the change to take effect. The default value is 0.0.0.0, although this is not a valid router ID.
LSA Type	The format and function of the link state advertisement: • **Router LSA**: A router can originate one or more Router LSAs for an area. A Router LSA originates in an area and describes the collected

Field	Description
	states of all the router's interfaces connected to the area. • **Network LSA**: A Network LSA is originated by the designated router for each link to which two or more routers are attached. The LSA lists all the routers attached to the link. • **Inter-Area Router LSA**: This type LSA is originated by an Area Border Router (ABR) and describes a prefix that is external to the area, yet internal to the autonomous system (AS). • **AS-External LSA**: This type of LSA is originated by an Autonomous System Border Router (ASBR) and describes a path to a prefix external to the autonomous system (AS). • **Link LSA**: A router originates a separate Link LSA for each attached link. It provides the router's link local address to routers that are attached to the link as well as a list of IPv6 prefixes that can be associated with the link. • **Intra-Area-Prefix LSA**: A link's designated router originates one or more Intra-Area-Prefix LSAs to advertise the link's prefixes throughout the area. A router can originate multiple Intra-Area-Prefix LSAs for an area to advertise its own prefixes and those of its attached stub links.
LS ID	The link state ID identifies the part of the routing domain that is being described by the advertisement. The value of the LS ID depends on the advertisement's LS type.
Age	The time in seconds since the link state advertisement was first originated.
Sequence	A signed 32-bit integer that is used to detect old and duplicate link state advertisements. The larger the sequence number, the more recent the advertisement.
Checksum	The checksum of the complete contents of the advertisement, except the LS age field. The checksum is used to detect data corruption of an advertisement. This corruption can occur while an advertisement is being flooded, or while it is being held in a router's memory.

Configure an OSPFv3 virtual link

If you cannot set up a physical link from the switch to the OSPFv3 backbone area, you can add a virtual link to connect to the backbone area through a common (non-backbone) area. You can configure a virtual link between any area border routers that have physical interfaces to the common area.

To configure an OSPFv3 virtual link:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Advanced > Virtual Link Configuration**.
   The OSPFv3 Virtual Link Configuration page displays.
3. In the **Area ID** field, enter an OSPFv3 area ID.
   An area ID is a 32-bit integer in dotted-decimal format that identifies the area to which a switch interface connects.
4. In the **Neighbor Router ID** field, enter the neighbor portion of a virtual link specification.
5. In the **Hello Interval** field, enter the OSPFv3 hello interval for the interface.
   This interval must be the same for all routers and switches in the network. The interval ranges from 1 to 65,535 seconds. The default is 10 seconds.
6. In the **Dead Interval** field, enter the number of seconds that the interface waits for a neighbor's hello packets before determining that the neighbor is down.
   This interval must be the same for all routers and switches in the network. This interval must be a multiple of the hello interval. The interval ranges from 1 to 65,535. The default is 40 seconds. (The default hello interval is 10 seconds.)
7. In the **Iftransit Delay Interval** field, enter the estimated number of seconds required to transmit a link state update packet over the interface.
   The interval range from 1 to 3600 seconds (1 hour). The default is 1 second.
8. In the **Retransmit Interval** field, enter the number of seconds between link-state advertisements for adjacencies, that is, for neighboring devices that form an OSPFv3 adjacency with the interface.
   The interval is also used when retransmitting database descriptions and link-state request packets. The value for the interval ranges from 1 to 3600 seconds (1 hour). The default is 5 seconds.
9. Click the **Add** button.
   Your settings are saved.

10. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
Neighbor State The state of the OSPFv3 neighbor interface:	
Down: This is the initial interface state. The interface is not usable for OSPFv3, and the interface settings are at their initial values. All interface timers are disabled, and no adjacencies are associated with the interface.	
Loopback: The interface is looped back either in hardware or software. The interface is unavailable for regular data traffic. You can get information about the interface by sending ICMP pings to the interface or through a bit error test because IP packets can still be addressed to an interface in a loopback state. An interface in a loopback state is advertised in LSAs as single host route for which the destination is the IP interface address.	
Waiting: The neighboring device is monitoring received hello packets to detect the identity of the backup designated router for the network. The neighboring device cannot elect a backup designated router or a designated router until it transitions out of the waiting state. This prevents unnecessary changes of the backup designated router.	
Designated Router: The neighboring device is the designated router on the attached network. Adjacencies are established to all other routers and switches that in the network. The neighboring device must also originate a network LSA for the network node. The network LSA contains links to all switches and routers (including the designated router) in the network.	
Backup Designated Router: The neighboring device is the backup designated router on the attached network. The neighboring device is promoted to designated router if the current designated router fails. The neighboring device establishes adjacencies to all other routers attached to the network. The backup designated router performs slightly different functions during the LSA flooding, as compared to the designated router.	
Other Designated Router: The interface is connected to a broadcast on which other routers are the designated router and backup designated router. The neighboring device attempts to form adjacencies to both the designated router and the backup designated router.	
State The state of the OSPFv3 switch interface:	
Down: This is the initial interface state. The interface is not usable for OSPFv3, and the interface settings are at their initial values. All interface timers are disabled, and no adjacencies are associated with the interface.	
Loopback: The interface is looped back either in hardware or software. The interface is unavailable for regular data traffic. You can get information about the interface by sending ICMP pings to the interface or through a bit error test because IP packets can still be addressed to an interface in a loopback state. An interface in a loopback state is advertised in LSAs as single host	

Field	Description
route for which the destination is the IP interface address.	
Waiting: The switch is monitoring received hello packets to detect the identity of the backup designated router for the network. The switch cannot elect a backup designated router or a designated router until it transitions out of the waiting state. This prevents unnecessary changes of the backup designated router.	
Designated Router: The switch is the designated router on the attached network. Adjacencies are established to all other routers and switches that in the network. The switch must also originate a network LSA for the network node. The network LSA contains links to all switches and routers (including the designated router) in the network.	
Backup Designated Router: The switch is the backup designated router on the attached network. The switch is promoted to designated router if the current designated router fails. The switch establishes adjacencies to all other routers attached to the network. The backup designated router performs slightly different functions during the LSA flooding, as compared to the designated router.	
Other Designated Router: The interface is connected to a broadcast on which other routers are the designated router and backup designated router. The switch attempts to form adjacencies to both the designated router and the backup designated router.	
Metric	The metric value that is used by the virtual link to compute the shortest paths.

Configure the OSPFv3 route redistribution

Route redistribution (RR) lets the switch share routing information between locally connected routers (or routing interfaces) and information about static routes.

To configure the OSPFv3 route redistribution:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Advanced > Route Redistribution**.
   The OSPFv3 Route Redistribution page displays.
3. Next to the **Source** field, select the check box for one or more types of source routes to which the configuration applies:
   - **Connected**: Routes between locally connected routers (or routing interfaces).
   - **Static**: Routes that were manually configured.

4. From the **Redistribute** menu, select if redistribution is enabled:
   - **Enable**: The type of route can be redistributed.
   - **Disable**: The type of route cannot be redistributed.
5. In the **Metric Value** field, set the metric value for the type of route.
   If the type of route is preconfigured, the metric value displays, but you can change it. The range is from 0 to 16777214.
6. From the **Metric Type** menu, select the type of metric that must used for the type of route:
   - **External Type 1**: External type 1 metrics that are comparable to the OSPFv3 metric.
   - **External Type 2**: External type 2 metrics that are assumed to be larger than the cost of the OSPFv3 metric.
7. In the **Tag** field, set the tag for the type of route.
   If the type of route is preconfigured, the tag displays, but you can change it. The range is from 0 to 4294967295.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Configure OSPFv3 nonstop forwarding

OSPFv3 nonstop forwarding (NSF) ensures that OSPF routing continues by temporarily letting other devices in the network perform OSPFv3 routing for the switch. That means that even if the switch restarts, OSPFv3 routing can continue. (We refer to this type of restart process as a graceful restart.)

To configure OSPFv3 NSF:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > OSPFv3 > Advanced > NSF OSPFv3 Summary**.
   The NSF OSPFv3 Summary page displays.
3. From the **Support Mode** menu, select how the switch performs a graceful restart:
   - **Always**: OSPFv3 performs a graceful restart for all planned and unplanned warm restart events.
   - **Disabled**: OSPFv3 graceful restarts are disabled. This is the default selection.
   - **Planned**: OSPFv3 performs a graceful restart only if a restart is planned (for example, because a CLI command is executed).

4. In the **Restart Interval** field, enter the time that the switch waits after a graceful restart before continuing OSPFv3 routing.
   The range is from 0 to 1800 seconds. The default is 120 seconds.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
Restart Status	The restart status of the OSPFv3 helper support mode: • Not Restarting • Planned Restart • Unplanned Restart
Restart Age (seconds)	The time in seconds since the last restart occurred.
Restart Exit Reason	Displays how the OSPFv3 routing process on the switch last started: • **Not Attempted**: A graceful restart was not attempted. • **Progress**: A graceful restart is in progress. • **Completed**: The previous graceful restart completed successfully. • **Timed Out**: The previous graceful restart timed out. • **Topology Changed**: The previous graceful restart terminated prematurely because of a topology change.

Configure Multicast Routing

This chapter covers the following topics:

- IPv4 multicast routing and the IPv4 multicast route table
- Distance Vector Multicast Routing Protocol
- IGMP for IPv4 multicast routing
- PIM for IPv4 multicast routing
- Static multicast routes for IPv4 addresses
- Multicast admin boundaries for IPv4 addresses
- IPv6 multicast routing and the IPv6 multicast route table
- PIM for IPv6 multicast routing
- MLD for IPv6 multicast routing
- Static multicast routes for IPv6 addresses

IPv4 multicast routing and the IPv4 multicast route table

Multicast is best suited for video and audio traffic requiring multicast packet control for optimal operation. Multicast for IPv4 includes support for IGMPv1, IGMPv2, and IGMPv3. Communication from point to multipoint is called multicasting. The source host (point) transmits a message to a group of zero or more hosts (multipoint) that are identified by a single IPv4 destination address. Although the task can be accomplished by sending unicast (point-to-point) messages to each of the destination hosts, multicasting is the preferred method for this type of transmission. A multicast message is delivered to all members of its destination host group with the same best-efforts reliability as regular unicast IPv4 messages. The message is not guaranteed to arrive intact at all members of the destination group or in the same order relative to other messages.

Display the IPv4 multicast route table

The multicast routing (Mroute) table includes information about the source IPv4 address, group destination IPv4 addresses, the incoming and outgoing interfaces, next hops to which IPv4 packet must be forwarded, and supported protocols.

To display the IPv4 Mroute table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > Mroute Table**.
   The Mroute Table page displays.

3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 137.Multicast IPv4 Mroute table information

Field	Description
Group IP	The destination group IPv4 address.
Source IP	The IPv4 address of the multicast packet source to be combined with the group IPv4 to fully identify a single route.
Incoming Interface	The incoming interface on which multicast packets for this source/group arrive.
Outgoing Interfaces	The list of outgoing interfaces on which multicast packets for this source/group are forwarded.
Up Time (hh:mm:ss)	The time in seconds since the entry was created.
Expiry Time (hh:mm:ss)	The time in seconds before this entry ages out and is removed from the table.
RPF Neighbor	The IPv4 address of the reverse path forwarding (RPF) neighbor.
Protocol	The multicast routing protocol which dynamically created this entry. The possible values are as follows: ○ PIM-DM ○ PIM-SM
Flags	The value displayed in this field is valid if the multicast routing protocol is PIM-SM. The possible values are RPT and SPT. For other protocols a "------" is displayed.

Add static multicast entries to the IPv4 Mroute table

You can add a static multicast entry to the IPv4 Mroute table.

To add a static multicast entry to the IPv4 Mroute table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > Mroute Static-Multicast Configuration**.
   The Mroute Static-Multicast Configuration page displays.
3. In the **Group IP** field, specify the multicast group IP address.

4. In the **Egress VLAN List** field, specify the VLAN or VLANs to which the multicast group IP address belongs.
5. Click the **Add** button.
   Your settings are saved. The multicast entry is added.
6. To refresh the page, click the **Refresh** button.
7. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 138.Mroute static-multicast configuration information

Field	Description
Maximum Multicast Static Address Count	The maximum number of static multicast addresses that the Mroute table can contain.
Current Multicast Static Address Count	The number of static multicast addresses that you added to the Mroute table.

Delete a static multicast entry from the IPv4 Mroute table

You can delete a static multicast entry that you no longer need from the IPv4 route table.

To delete a static multicast entry from the IPv4 Mroute table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > Mroute Static-Multicast Configuration**.
   The Mroute Static-Multicast Configuration page displays.
3. Select the check box for the multicast entry.
4. Click the **Delete** button.
   Your settings are saved. The multicast entry is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure global multicast settings for the switch

You can configure the global multicast settings for the switch, as opposed to the multicast settings for an interface.

To configure global multicast settings for the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > Global Configuration**.
   The Global Configuration page displays.
3. Select the Admin Mode **Enable** or **Disable** radio button to enable or disable the administrative status of multicast forwarding on the switch.
   The default is Disable.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 139.Multicast global configuration information

Field	Description
Protocol State	The operational state of multicast forwarding.
Table Maximum Entry Count	The maximum number of entries in the Mroute table.
Protocol	The multicast routing protocol that is active on the switch, if any protocol is active.
Table Entry Count	The number of multicast route entries present in the Mroute table.

Configure a multicast interface

You can configure one or more multicast interfaces.

To configure a multicast interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > Interface Configuration**.
   The Interface Configuration page displays.

3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. In the **TTL Threshold** field, specify the threshold below which a multicast data packet is not forwarded from the interface.
   Enter a number between 0 and 255. The default is 1. For a packet to be forwarded, its multicast time-to-live (TTL) value must be greater than the value that you set. If you enter 0, all multicast packets on the interface are forwarded.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Distance Vector Multicast Routing Protocol

The Distance Vector Multicast Routing Protocol (DVMRP) is used for multicasting over IP networks without routing protocols to support multicast. The DVMRP is based on the RIP protocol but more complex than RIP.
The DVMRP operates as follows:

1. The first message (such as an IGMP message) for any source-group pair is forwarded to the entire multicast network, with respect to the time-to-live (TTL) of the packet.
2. TTL restricts the area to be flooded by the message.
3. All the leaf routers that do not include members on directly attached subnetworks send back prune messages to the upstream router.
4. The branch that transmits a prune message is deleted from the delivery tree.

5. The delivery tree, which is spanning to all the members in the multicast group, is constructed.

In this way, DVRMP can maintain a link-state database that keeps track of the return paths to the source of multicast packages.

Enable DVMRP on the switch and view route information

You can enable DVMRP on the switch and view route information.

To enable DVMRP on the switch and view global route information:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > DVMRP > Global Configuration**.
   The DVMRP Global Configuration page displays.
3. Select an Admin Mode radio button:
   - **Enable**: DVMRP is enabled for the switch.
   - **Disable**: DVMRP is disabled for the switch. This is the default setting.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
Version	The protocol version of the DVMRP, which is always 3.
Total Number of Routes	The number of routes in the DVMRP routing table.
Reachable Routes	The number of reachable routes in the DVMRP routing table. A reachable route has a metric that lets the switch reach the endpoint.

Configure a DVMRP interface

You can configure one or more DVMRP interfaces, which are interfaces from which the switch sends DVMRP messages that let the multicast delivery tree compute the route cost to reach the switch.

To configure a DVMRP interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > DVMRP > Interface Configuration**.
   The DVMRP Interface Configuration page displays.
3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces or VLANs by taking one of the following actions:
   - To configure a single interface or VLAN, select the check box associated with the interface or VLAN, or type the interface or VLAN number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces or VLANs with the same settings, select the check box associated with each interface or VLAN.
5. From the **Interface Mode** menu, select if DVMRP is enabled on the interface:
   - **Enable**: DVMRP is enabled for the interface.
   - **Disable**: DVMRP is disabled for the interface. This is the default setting.
6. In the **Interface Metric** field, enter the DVMRP metric for the interface.
   The interface sends the metric in DVMRP messages that state the route cost to reach the switch. The range is from 1 to 31. The default is 1.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Field	Description
Protocol State	The operational state of the DVMRP protocol on the interface: Operational or Non-Operational.
Local Address	The IP address that is used as a source address in packets sent from the interface.
Generation ID	The DVMRP generation ID used by the switch for the interface. This ID is reset each time the interface is started and is sent in prune messages. A change in generation ID informs the neighbor routers to discard any previous information about the switch.
Received Bad Packets	The number of invalid packets received on the interface.
Received Bad Routes	The number of invalid routes received on the interface.
Sent Routes	The number of routes sent on the interface.

View DVMRP neighbors

If DVMRP is enabled, you can view DVMRP neighbors.

To view the DVMRP neighbors:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > DVMRP > DVMRP Neighbor**.
   The DVMRP Neighbor page displays.
3. To search for a neighbor, do the following:
   a. From the **Search** menu, select **Interface** or **Neighbor IP**.
   b. In the **Search** field, enter an interface number in the unit/slot/port format (for example, 1/0/14) or IP address, depending on your selection from the menu.
   c. Click the **Go** button.
4. To refresh the page, click the **Refresh** button.

The following tables describe the view-only fields on the page.

Field	Description
Interface	The interface for which data is displayed.
Neighbor IP	The IP address of the neighbor as detected by the interface.
State	The state of the neighbor router as detected by the interface: active or down.
Up Time	The DVMRP uptime for the neighbor as detected by the interface. This is the period since the neighbor entry was detected by the interface.
Expiry Time	The DVMRP expiration time for the neighbor as detected by the interface. This is the period that is left before the neighbor entry ages out, and is not applicable if the neighbor router's state is down.
Generation ID	The DVMRP generation ID for the neighbor as detected by the interface.
Major Version	The DVMRP major version for the neighbor as detected by the interface. For example, if the neighbour is running version 3 of DVMRP, the major version is set as 0x03 and the minor version is set as 0xFF.
Minor Version	The DVMRP minor version for the neighbor as detected by the interface.
Capabilities	The DVMRP capabilities of the neighbor as detected by the interface.
Received Routes	The number of routes received for the neighbor as detected by the interface.
Received Bad Packets	The number of invalid packets received for the neighbor as detected by the interface.
Received Bad Routes	The number of invalid routes received for the neighbor as detected by the interface.

View the DVMRP next hops

If DVMRP is enabled, you can view DVMRP the next hops.

To view the DVMRP next hops:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Routing > Multicast > DVMRP > DVMRP Next Hop**.
   The DVMRP Next Hop page displays.

> **NOTE:** If you select to display all pages instead of a limited number, it might take a while before the information is displayed.

3. To refresh the page, click the **Refresh** button.
   The following tables describe the view-only fields on the page.

Field	Description
Source IP	The IP address that is used with the source mask to identify the source network
Source Mask	The network mask used with the source IP address
Next Hop Interface	The next hop's interface that connects to the switch's interface
Type	The next hop type: Leaf means that no downstream dependent neighbors exist on the outgoing interface. Otherwise, the type is branch. • **Leaf**: No downstream dependent neighbors exist on the next hop interface • **Branch**: Downstream dependent neighbors do exist on the next hop interface

View the DVMRP prune table

If DVMRP is enabled, you can view DVMRP prune table, which shows the multicast group IP addresses that were removed from the multicast delivery tree because they no longer serve any multicast members.

To view the DVMRP prune table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > DVMRP > DVMRP Prune**.
   The DVMRP Prune page displays.
3. To refresh the page, click the **Refresh** button.

The following tables describe the view-only fields on the page.

Field	Description
Group IP	The group IP address that was pruned.
Source IP	The IP address that is used with the source mask to identify the source network.
Source Mask	The network mask used with the source IP address.
Expiry Time	The period remaining before the prune expires at the upstream neighbor. If no prune messages are received from downstream neighbors, this period is set to the value of the default prune lifetime time; otherwise, it is set to the smallest received value or the default timer, whichever is smaller.

View the DVMRP routes

If DVMRP is enabled, you can view DVMRP the routes that are being advertised.

To view the DVMRP routes:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > DVMRP > DVMRP Route**.
   The DVMRP Route page displays.
3. To refresh the page, click the **Refresh** button.

The following tables describe the view-only fields on the page.

Table 140.DVMRP route information

Field	Description
Source IP	The IP address that is used with the source mask to identify the source network.
Source Mask	The network mask used with the source IP address.
Upstream Neighbor	The address of the upstream neighbor (for example, a reverse-path forwarding [RPF] neighbor) from which IP datagrams are received.
Interface	The switch interface on which IP datagrams from the upstream neighbors are received. A value of 0 means that the route is an aggregate for which no next-hop interface exists.
Metric	The distance in hops to the source subnet.
Expiry Time	The period remaining before the route entry expires.
Up Time	The period since the route was detected by the switch.

IGMP for IPv4 multicast routing

You can configure the Internet Group Management Protocol (IGMP) settings and view the IGMP statistics.
IGMP snooping lets the switch forward multicast traffic intelligently. Based on the IGMP query and report messages, the switch forwards traffic only to the ports that request the multicast traffic, rather than to all ports.
IGMP snooping helps to optimize multicast performance and is especially useful for bandwidth-intensive IP multicast applications such as online media streaming.

Enable or disable IGMP for the switch

You can globally enable or disable IGMP for the switch.

To configure global IGMP settings for the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > IGMP > Global Configuration**.
   The Global Configuration page displays.

3. Select the Admin Mode **Enable** or **Disable** radio button to enable or disable IGMP on the switch.
   The default is Disable.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure an IGMP routing interface

You can configure one or more IGMP routing interfaces.

To configure an IGMP routing interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select the **Routing > Multicast > IGMP > Routing Interface Configuration**.
   The Routing Interface Configuration page displays.
3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Admin Mode** menu, select **Enable** or **Disable** to enable or disable IGMP for the interface.
   The default is Disable.

6. In the **Version** field, specify the version of IGMP for the interface.
   The version can be 1, 2, or 3. The default is 3. This field is configurable only when IGMP Interface mode is enabled.
7. In the **Robustness** field, enter the robustness value for the interface.
   The robustness value allows tuning to compensate for the expected packet loss on a subnet. IGMP is robust to (robustness value - 1) packet losses. If you expect the subnet to be lossy, enter a higher number. The value is from 1 to 255. The default is 2.
8. In the **Query Interval** field, enter the frequency in seconds at which IGMP host membership query packets are transmitted on the interface.
   The value is from 1 to 3600. The default is 125.
9. In the **Query Max Response Time** field, enter the maximum query response time, in tenths of a second, for a Membership Report to be received before the IGMPv2 query expires on the interface.
   The value is from 0 to 255. The default is 100.
10. In the **Startup Query Interval** field, enter the number of seconds between the transmission of startup queries on the interface.
    The value is from 1 to 300. The default is 31.
11. In the **Startup Query Count** field, enter the number of queries to be sent on startup, with the time between queries determined by the startup query interval.
    The value is from 1 to 20. The default is 2.
12. In the **Last Member Query Interval** field, enter the last member query interval in tenths of a second.
    This value is the maximum response time to be inserted into group-specific queries sent in response to leave group messages, and is also the period between group-specific query messages. The value is from 0 to 255. The default is 10. This value is not used for IGMP version 1.
13. In the **Last Member Query Count** field, enter the number of queries to be sent on receiving a leave group report.
    The value is from 1 to 20. The default is 2.
14. Click the **Apply** button.
    Your settings are saved.
15. To save the settings to the running configuration, click the **Save** icon.

Display the statistics for the IGMP routing interfaces

You can display the statics for the IGMP routing interfaces.

To display the statistics for the IGMP routing interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > IGMP > Routing Interface Statistics**.
   The Routing Interface Statistic page displays.
3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 141.IGMP routing interface statistics information

Field	Description
Interface	The interface on which IGMP is enabled.
IP Address	The IP address of the interface.
Subnet Mask	The subnet mask for the IP address of the interface.
Protocol State	The operational state of IGMP on the interface (Operational or Non-Operational).
Querier IP	The address of the IGMP querier on the IP subnet to which the interface is attached.
Querier Status	Indicates if the interface is in querier or non-querier mode.
Querier Up Time	The period in seconds since the IGMP interface querier was last changed.
Querier Expiry Time	The period in seconds remaining before the other querier present timer expires. If the switch is the querier, this is zero.
Wrong Version Queries Received	The number of queries that were received on the interface with an IGMP version that does not match the IGMP version configured for the interface, over the lifetime of the entry. IGMP requires that all routers or routing interfaces on a LAN are configured to run the same version of IGMP. Therefore, a configuration error is indicated if any queries are received with the wrong version number.
Number of Joins Received	The number of times a group membership was added on the interface; that is, the number of times an entry for the interface was added to the cache table. This gives an indication of the amount of IGMP activity on the interface.
Number of Groups	The current number of entries for the interface in the cache table.

Display the IGMP groups and search the IGMP group database

You can display the IGMP group information and search the IGMP group database by interface or by group.

To display the IGMP groups and search the IGMP group database:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > IGMP > IGMP Groups**.
   The IGMP Groups page displays.
3. To search for IGMP group entries by interface or group, select one of the following options from the **Search** menu:
   - **Interface**: Select **Interface** from the menu, specify the interface in unit/port format (for example 0/8), and click the **Go** button. If the entry exists, the entry is displayed as the first entry, followed by the remaining entries.
   - **Group**: Select **Group** from the menu, specify the multicast group IPv4 address, and click the **Go** button. If the entry exists, that entry with the matching group IPv4 address is displayed as the first entry, followed by the remaining entries. An exact match is required.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 142.IGMP group information

Field	Description
Interface	The interface on which IGMP is enabled.
Multicast Group IP	The IP multicast group address.
Last Reporter	The IP address of the source of the last membership report that was received for the IP multicast group address on the interface.
Up Time	The period elapsed since this entry was created.
Expiry Time	The period remaining before this entry is aged out.
Version 1 Host Timer	The period remaining until the switch determines that no IGMP version 1 members are on the IP subnet attached to the interface. When an IGMPv1 membership report is received, this timer is reset to the group membership timer. Unless this timer reaches zero, the switch ignores any IGMPv2 leave messages for this group that it receives on the interface. This field displays only if the interface is configured for IGMP version 1.
Version 2 Host Timer	The time remaining until the switch determines that no IGMP version 2 members are on the IP subnet attached to the interface. When an IGMPv2 membership report is received, this timer is reset to the group membership timer. Unless this timer reaches zero, the switch ignores any IGMPv1 and IGMPv3 leave messages for this group that it receives on the selected interface. This field displays only if the interface is configured for IGMP version 2.
Compatibility	The group compatibility mode (v1, v2, or v3) for the group on the interface.
Filter Mode	The source filter mode (Include, Exclude, or NA) for the group on the interface. If the mode is NA, the field is blank.

Display the IGMP membership information and search the IGMP membership database

You can display the IGMP membership information and search the IGMP membership database by interface or by group.

To display the IGMP membership information and search the IGMP membership database:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > IGMP > IGMP Membership**.
   The IGMP Membership page displays.
3. To search for IGMP membership by interface or group, select one of the following options from the **Search** menu:
   - **Interface**: Select **Interface** from the menu, specify the interface in unit/port format (for example 0/8), and click the **Go** button. If the entry exists, the entry is displayed as the first entry, followed by the remaining entries.
   - **Group**: Select **Group** from the menu, specify the multicast group IP address, and click the **Go** button. If the entry exists, that entry with the matching group IP address is displayed as the first entry, followed by the remaining entries. An exact match is required.
4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 143.IGMP membership information

Field	Description
Interface	The interface on which multicast packets are forwarded.
Group IP	The IP multicast group address.
Compatibility Mode	The group compatibility mode (v1, v2, or v3) for the group on the interface.
Source Filter Mode	The source filter mode (Include, Exclude, or NA) for the group on the interface. If the mode is NA, the field is blank.
Source Hosts	The source IP addresses that are members of the multicast address.
Expiry Time	The expiration time that applies to each source IP address that is a member of the multicast group. This is the period after which the source entry ages out.

Configure an IGMP proxy interface

You can configure an IGMP proxy interface. The following requirements apply:

- Routing, IGMP, and multicast must be globally enabled on the switch.
- You must configure at least one routing interface.
- The IGMP proxy interface cannot be the same interface as an IGMP routing interface.

To configure an IGMP proxy interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > IGMP > Proxy Interface Configuration**.
   The Proxy Interface Configuration page displays.
3. From the **Interface** menu, select the interface that must function as the IGMP proxy interface.
4. From the **Admin Mode** menu, select **Enable** or **Disable** to enable or disable the IGMP proxy capability on the interface.
   The default is Disable.
5. In the **Unsolicited Report Interval** field, specify the period in seconds between repetitions of a host's initial report of membership in a group.
   The range is from 1 to 260. The default is 1.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 144.IGMP proxy interface configuration information

Field	Description
IP Address	The IP address of the IGMP proxy interface.
Subnet Mask	The subnet mask for the IP address of the IGMP proxy interface.
Operational Mode	The operational state of IGMP proxy interface (Disable or Enable).
Querier Address on Proxy Interface	The querier address on the proxy interface.
Number of Groups	The number of multicast group entries for the IGMP proxy interface in the cache table.
Version	The version of IGMP that is configured on the IGMP routing interface that you changed into the IGMP proxy interface. For more information, see Configure an IGMP routing interface on page 477.
Version 1 Querier Timeout	The IGMPv1 querier time-out period in seconds. If a host receives an IGMPv1 query, the host temporarily sets its IGMPv1 querier present timer to the IGMPv1 querier Interval. After the host reads the IGMPv1 query, the IGMPv1 querier time-out indicates the period after which the host transitions back to IGMPv3.
Version 2 Querier Timeout	The IGMPv2 querier time-out period in seconds. If a host receives an IGMPv2 query, the host temporarily sets its IGMPv2 querier present timer to the IGMPv2 querier Interval. After the host reads the IGMPv2 query, the IGMPv2 querier time-out indicates the period after which the host transitions back to IGMPv3.
Proxy Start Frequency	The number of times the proxy was started.

Display the statistics for the IGMP proxy interface

You can display the statistics for the IGMP proxy interface.

To display the statistics for the IGMP proxy interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Routing > Multicast > IGMP > Proxy Interface Statistics**.
   The Proxy Interface Statistics page displays.
3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 145.IGMP proxy interface statistics information

Field	Description
Proxy Interface	The interface on which IGMP packets are received.
Version	The version of IGMP packets received.
Queries Received	The number of IGMP queries received.
Report Received	The number of IGMP reports received.
Reports Sent	The number of IGMP reports sent.
Leaves Received	The number of IGMP leave messages received.
Leaves Sent	The number of IGMP leave messages sent.

Display the IGMP proxy membership and search the IGMP proxy membership database

You can display the IGMP proxy membership information and search the IGMP proxy membership database by group.

To display the IGMP proxy membership information and search the IGMP proxy membership database:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > IGMP > Proxy Membership**.
   The Proxy Membership page displays.
3. To search for IGMP proxy membership by group, specify the multicast group IPv4 address, and click the **Go** button.
   If the entry exists, that entry with the matching group IPv4 address is displayed as the first entry, followed by the remaining entries. An exact match is required.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 146.IGMP proxy membership information

Field	Description
Proxy Interface	The interface on which the IGMP proxy is enabled.
Group IP	The multicast group IPv4 address.
Source Hosts	The source addresses that are members of the multicast group.
Last Reporter	The source IPv4 address that sent the last membership report for the multicast group on the IGMP proxy interface.
Up Time	The time elapsed since the entry was created.
Expiry Time	The expiration interval for each source address that is a member of the multicast group. This is the period after which the source entry ages out.
State	The state of the host entry: ○ **Non-member state**: The host does not belong to the group on the interface. ○ **Delaying member state**: The host belongs to the group on the interface and its report timer is active. (The report timer is used to send reports.) ○ **Idle member state**: The host belongs to the group on the interface but its report timer is inactive.
Filter Mode	The group filter mode (Include, Exclude, or None) for the group on the IGMP proxy interface.
Number of Sources	The number of source hosts in the multicast group.

PIM for IPv4 multicast routing

Protocol-Independent Multicast (PIM) is a standard multicast routing protocol that provides scalable interdomain multicast routing across the Internet, independent of any particular unicast routing protocol.
You can configure the various PIM settings for IPv4 multicast routing and display the PIM statistics.

Configure the global PIM IPv4 settings on the switch

Depending on your IPv4 network needs, you can globally enable PIM Sparse Mode (PIM-SM) or PIM Dense Mode (PIM-DM) on the switch.

To configure the global PIM IPv4 settings on the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > PIM > Global Configuration**.
   The Global Configuration page displays.
3. Select one of the following Admin Mode radio buttons:
   - **Disable**: PIM is disabled on the switch. This is the default setting.
   - **PIM-SM**: PIM Sparse Mode is enabled on the switch.
   - **PIM-DM**: PIM Dense Mode is enabled on the switch.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Add IPv4 PIM-SSM groups

While PIM employs a specially configured rendezvous point (RP) router that serves as a meeting junction for multicast senders and listeners, PIM Single-Source Multicast (PIM-SSM) does not use a rendezvous point (RP). PIM-SSM supports source route delivery trees only. It is used between routers (including switch router interfaces) so that they can track which multicast packets to forward to each other and to their directly connected LANs. You can implement PIM-SSM with a strict subset of the PIM protocol mechanisms. Both regular IPv4 multicast and SSM can coexist on a single router, and both can be implemented using the PIM protocol. A range of multicast addresses, 232.0.0.0/8 for IPv4, is reserved for SSM.

To add an IPv4 PIM-SSM group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > PIM > SSM Configuration**.
   The SSM Configuration page displays.
3. In the **SSM Group Address** field, specify the source-specific multicast group IPv4 address.

4. In the **SSM Group Mask** field, specify the mask for the source-specific multicast group IPv4 address.
5. Click the **Add** button.
   Your settings are saved. The group is added.
6. To save the settings to the running configuration, click the **Save** icon.

Delete an IPv4 PIM-SSM group

You can delete an IPv4 PIM-SSM group that you no longer need.

To delete an IPv4 PIM-SSM group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > Mroute Static-Multicast Configuration**.
   The Mroute Static-Multicast Configuration page displays.
3. Select the check box for the PIM-SSM group.
4. Click the **Delete** button.
   Your settings are saved. The PIM-SSM group is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure an IPv4 PIM interface

You can configure an IPv4 interface for PIM.

To configure an IPv4 PIM interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > PIM > Interface Configuration**.
   The Interface Configuration page displays.

3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Admin Mode** menu, select **Enable** or **Disable** to enable or disable PIM for the interface.
   The default is Disable.
6. In the **Hello Interval** field, specify the period in seconds between the transmission of PIM hello messages on the interface.
   The range is from 0 to 18000. The default is 30.
7. In the **Join/Prune Interval** field, specify period in seconds at which PIM Join/Prune messages are transmitted on the interface.
   The range is from 0 to 18000. The default is 60.
8. From the **BSR Border** menu, select **Enable** or **Disable** to enable or disable the bootstrap router (BSR) border on the interface.
   The default is Disable.
9. In the **DR Priority** field, specify the designated router (DR) priority for the interface.
   The range is from 0 to 2147483647. The default is 1.
10. Click the **Apply** button.
    Your settings are saved.
11. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only field on the page.

Table 147.IPv4 PIM interface configuration information

Field	Description
Protocol State	The state of PIM on the interface (Operational or Non-Operational).
IP Address	The IPv4 address of the PIM interface.
Designated Router	The IPv4 address of the designated router on the PIM interface.
Neighbor Count	The number of PIM neighbors on the PIM interface.

Display IPv4 PIM neighbors and search the PIM neighbor database

You can display the IPv4 PIM neighbor information and search the PIM neighbors database by interface or by neighbor IP address.

To display IPv4 PIM neighbors and search the PIM neighbors database:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > PIM > PIM Neighbor**.
   The PIM Neighbor page displays.
3. To search for IGMP group entries by interface or group, select one of the following options from the **Search** menu:
   - **Interface**: Select **Interface** from the menu, specify the interface in unit/port format (for example 0/8), and click the **Go** button. If the entry exists, the entry is displayed as the first entry, followed by the remaining entries.
   - **Neighbor IP**: Select **Neighbor IP** from the menu, specify the neighbor IPv4 address, and click the **Go** button. If the entry exists, that entry with the matching neighbor IPv4 address is displayed as the first entry, followed by the remaining entries. An exact match is required.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 148.IPv4 PIM neighbor information

Field	Description
Interface	The interface on which the PIM neighbor is connected.
Neighbor IP	The IPv4 address of the PIM neighbor.
Up Time (hh:mm:ss)	The time that passed since the PIM device became a neighbor of the interface.
Expiry Time (hh:mm:ss)	The time remaining before the PIM neighbor ages out.
DR Priority	The designated router (DR) priority of the PIM neighbor.

Add an IPv4 PIM candidate rendezvous point configuration

You can add an IPv4 PIM candidate rendezvous point (RP) configuration on an interface.

To add an IPv4 PIM candidate RP configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > PIM > Candidate RP Configuration**.
   The Candidate RP Configuration page displays.
3. From the **Interface** menu, select the interface.
4. In the **Group Address** field, specify the group IPv4 address that is transmitted in candidate-RP-advertisements.
5. In the **Group Mask** field, enter the group address mask that is transmitted in candidate-RP-advertisements.
6. In the **C-RP Advertisement Interval** field, specify the period in seconds at which candidate-RP-advertisements are sent as unicast traffic to the bootstrap router (BSR).
   The range is from 1 to 16383 seconds.
7. Click the **Add** button.
   Your settings are saved. The PIM candidate RP is added.
8. To save the settings to the running configuration, click the **Save** icon.

Delete an IPv4 PIM candidate rendezvous point configuration

You can delete an IPv4 PIM candidate rendezvous point (RP) configuration that you no longer need.

To delete an IPv4 PIM candidate RP configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > PIM > Candidate RP Configuration**.
   The Candidate RP Configuration page displays.
3. From the **Interface** menu, select the interface.
4. Select the check box for the PIM candidate RP.
5. Click the **Delete** button.
   Your settings are saved. The PIM candidate RP entry is removed.

Configure an interface as an IPv4 PIM bootstrap router candidate

You can configure an interface as an IPv4 PIM bootstrap router (BSR) candidate.

To configure an interface as an IPv4 PIM BSR candidate:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > PIM > BSR Candidate Configuration**.
   The BSR Candidate Configuration page displays.
3. From the **Interface** menu, select an interface.
4. In the **Hash Mask Length** field, specify the hash mask length to be advertised in bootstrap messages.
   The hash mask length is used in the hash algorithm for selecting the RP for a particular group. The range is from 0 to 32. The default is 30.
5. In the **Priority** field, specify the priority for the BSR candidate.
   The range is from 0 to 255. The default is 0.
6. In the **Advertisement Interval** field, specify the period in seconds between advertisements.
   The range is from 1 to 16383. The default is 60.

7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.
9. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 149.IPv4 PIM BSR candidate configuration information

Field	Description
BSR Expiry Time (hh:mm:ss)	The time (in hours, minutes and seconds) after which the learned elected BSR expires.
IP Address	The IPv4 address of the elected BSR.
Next bootstrap Message (hh:mm:ss)	The time (in hours, minutes, and seconds) after which the next bootstrap message is due from the BSR.
Next Candidate RP Advertisement (hh:mm:ss)	The time (in hours, minutes, and seconds) after which the next candidate RP advertisement is sent.

Delete an IPv4 PIM bootstrap router candidate configuration

You can delete an IPv4 PIM bootstrap router (BSR) candidate configuration from an interface.

To delete an IPv4 PIM BSR candidate configuration from an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > PIM > BSR Candidate Configuration**.
   The BSR Candidate Configuration page displays.
3. From the **Interface** menu, select the interface.
4. Click the **Delete** button.
   Your settings are saved. The PIM BSR candidate configuration is deleted from the interface.
5. To save the settings to the running configuration, click the **Save** icon.

Configure a static IPv4 PIM rendezvous point for a group

You can add a new or change an existing static IPv4 PIM rendezvous point (RP) for a group.

To add or change a static IPv4 PIM RP for a group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > PIM > Static RP Configuration**.
   The Static RP Configuration page displays.
3. If you are changing an existing static RP, select the check box for the static RP.
4. In the **RP Address** field, specify the IPv4 address for the static RP.
5. In the **Group Address** field, specify the IPv4 address for the group to which the static RP applies.
6. In the **Group Mask** field, specify the mask for the group to which the static RP applies.
7. From the **Override** menu, select **Enable** or **Disable** to enable or disable the override capacity of a static RP over an RP that is learned from a BSR.
   If you select **Enable** and conflict occurs, a static RP overrides an RP that is learned from a BSR.
8. Do one of the following:
   - If you are adding a new static RP, click the **Add** button. Your settings are saved. The static RP is added.
   - If you are changing an existing static RP, click the **Apply** button. Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Delete a static IPv4 PIM rendezvous point configuration

You can delete a static IPv4 PIM rendezvous point (RP) configuration that you no longer need for a group.

To delete a static IPv4 PIM RP configuration for a group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > PIM > Static RP Configuration.**
   The Static RP Configuration page displays.
3. Select the check box for the static RP.
4. Click the **Delete** button.
   Your settings are saved. The static RP configuration is deleted for the group.
5. To save the settings to the running configuration, click the **Save** icon.

Static multicast routes for IPv4 addresses

If a multicast routing protocol such as PIM-SM or PIM-DM is enabled and configured, multicast routes are built dynamically. However, you can also configure one or more static IPv4 multicast routes. A static route can facilitate the processing of multicast traffic when a neighboring reverse-path forwarding (RPF) router is located between the switch and the multicast source.

Configure static multicast routes for IPv4 addresses

You can add a new or change an existing static IPv4 multicast route.

To add or change a static IPv4 multicast route:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > Static Routes Configuration**.
   The Static Routes Configuration page displays.
3. If you are changing an existing static multicast route, select the check box for the route.
   For an existing static route, you can change the metric only.
4. In the **Source IP** field, specify the IPv4 address of the multicast packet source.
5. In the **Source Mask** field, specify the subnet mask of the multicast packet source.
6. In **RPF Neighbor** field, specify the IPv4 address of the neighbor RPF router that is located between the switch and the multicast source.
7. In the **Metric** field, specify the link state cost of the path from the multicast source to the switch.
   The range is from 0 to 255.
8. Do one of the following:
   - If you are adding a new static multicast route, click the **Add** button. Your settings are saved. The static multicast route is added.
   - If you are changing an existing static multicast route, click the **Apply** button. Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Delete a static multicast route for an IPv4 address

You can delete a static IPV4 multicast route that you no longer need.

To delete a static IPv4 multicast route:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > Static Routes Configuration**.
   The Static Routes Configuration page displays.
3. Select the check box for the static route.
4. Click the **Delete** button.
   Your settings are saved. The static route is deleted.
5. To save the settings to the running configuration, click the **Save** icon.

Multicast admin boundaries for IPv4 addresses

A multicast administrative (admin) boundary can stop ingress and egress multicast traffic for a range of multicast addresses on routing interface. Using admin boundaries, you can set the borders for a multicast region, beyond which multicast traffic is not supposed to cross.

Configure an interface as a multicast admin boundary

You can configure an interface as a multicast admin boundary.

To configure an interface as a multicast admin boundary:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > Admin Boundary Configuration**.
   The Admin Boundary Configuration page displays.
3. From the **Interface** menu, select the router interface.
4. In the **Group IP** field, specify the multicast group address for the start of the range of addresses that must be excluded.
   The address must be in the range from 239.0.0.0 through 239.255.255.255.
5. In the **Group Mask** field, enter the mask for the multicast group address.
   The combination of the mask and the group IP defines the range of the multicast admin boundary.
6. Click the **Add** button.
   Your settings are saved. The multicast admin boundary is added.
7. To save the settings to the running configuration, click the **Save** icon.

Delete a multicast admin boundary configuration for an interface

You can delete a multicast admin boundary configuration that you no longer need for an interface.

To delete a multicast admin boundary configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > Multicast > Admin Boundary Configuration**.
   The Admin Boundary Configuration page displays.
3. Select the check box for the interface.
4. Click the **Delete** button.
   Your settings are saved. The static route is deleted from the interface.
5. To save the settings to the running configuration, click the **Save** icon.

IPv6 multicast routing and the IPv6 multicast route table

Multicast is best suited for video and audio traffic requiring multicast packet control for optimal operation. Multicast for IPv6 includes support for Multicast Listener Discovery (MLD). Communication from point to multipoint is called multicasting. The source host (point) transmits a message to a group of zero or more hosts (multipoint) that are identified by a single IPv6 destination address. Although the task can be accomplished by sending unicast (point-to-point) messages to each of the destination hosts, multicasting is the preferred method for this type of transmission. A multicast message is delivered to all members of its destination host group with the same best-efforts reliability as regular unicast IPv6 messages. The message is not guaranteed to arrive intact at all members of the destination group or in the same order relative to other messages.

Display the IPv6 multicast route table

The multicast routing (Mroute) table includes information about the source IPv6 address, group destination IPv6 addresses, the incoming and outgoing interfaces, next hops to which IPv6 packet must be forwarded, and supported protocols.

To display the IPv6 Mroute table:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Routing > IPv6 Multicast > Mroute Table**.
   The Mroute Table page displays.
3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 150.Multicast IPv6 Mroute table information

Field	Description
Group IP	The destination group IPv6 address.
Source IP	The IPv6 address of the multicast packet source to be combined with the group IPv6 to fully identify the single route.
Incoming Interface	The incoming interface on which multicast packets for this source/group arrive.
Outgoing Interfaces	The list of outgoing interfaces on which multicast packets for this source/group are forwarded.
Up Time (hh:mm:ss)	The time in seconds since the entry was created.
Expiry Time (hh:mm:ss)	The time in seconds before this entry ages out and is removed from the table.
RPF Neighbor	The IPv6 address of the reverse path forwarding (RPF) neighbor.
Protocol	The multicast routing protocol which dynamically created this entry. The possible values are as follows: ○ PIM-DM ○ PIM-SM
Flags	The value displayed in this field is valid if the multicast routing protocol is PIM-SM. The possible values are RPT and SPT. For other protocols a "------" is displayed.

PIM for IPv6 multicast routing

Protocol-Independent Multicast (PIM) is a standard multicast routing protocol that provides scalable interdomain multicast routing across the Internet, independent any particular unicast routing protocol.
You can configure the various PIM settings for IPv6 multicast routing and display the PIM statistics.

Configure the global PIM IPv6 settings on the switch

Depending on your IPv6 network needs, you can globally enable PIM Sparse Mode (PIM-SM) or PIM Dense Mode (PIM-DM) on the switch.

To configure the global PIM IPv6 settings on the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > IPv6 PIM > Global Configuration**.
   The Global Configuration page displays.
3. Select one of the following Admin Mode radio buttons:
   - **Disable**: PIM is disabled on the switch. This is the default setting.
   - **PIM-SM**: PIM Sparse Mode is enabled on the switch.
   - **PIM-DM**: PIM Dense Mode is enabled on the switch.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Add IPv6 PIM-SSM groups

While PIM employs a specially configured rendezvous point (RP) router that serves as a meeting junction for multicast senders and listeners, PIM Single-Source Multicast (PIM-SSM) does not use a rendezvous point (RP). PIM-SSM supports source route delivery trees only. It is used between routers (including switch router interfaces) so that they can track which multicast packets to forward to each other and to their directly connected LANs. You can implement PIM-SSM with a strict subset of the PIM protocol mechanisms. Both regular IPv6 multicast and SSM can coexist on a single router, and both can be implemented using the PIM protocol. A range of multicast addresses, FF3x::/32 for IPv6, is reserved for SSM.

To add an IPv6 PIM-SSM group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > IPv6 PIM > SSM Configuration**.
   The SSM Configuration page displays.
3. In the **SSM Group Address** field, enter the source-specific multicast group IPv6 address.
4. In the **SSM Group Mask** field, enter the source-specific multicast group IPv6 address prefix.

5. Click the **Add** button.
   Your settings are saved. The group is added.
6. To save the settings to the running configuration, click the **Save** icon.

Delete an IPv6 PIM-SSM group

You can delete an IPv6 PIM-SSM group that you no longer need.

To delete an IPv6 PIM-SSM group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > IPv6 PIM > SSM Configuration**.
   The SSM Configuration page displays.
3. Select the check box for the PIM-SSM group.
4. Click the **Delete** button.
   Your settings are saved. The PIM-SSM group is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure an IPv6 PIM interface

You can configure an IPv6 interface for PIM.

To configure an IPv6 PIM interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > IPv6 PIM > Interface Configuration**.
   The Interface Configuration page displays.

3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Admin Mode** menu, select **Enable** or **Disable** to enable or disable PIM for the interface.
   The default is Disable.
6. In the **Hello Interval** field, specify the period in seconds between the transmission of PIM hello messages on the interface.
   The range is from 0 to 18000. The default is 30.
7. In the **Join/Prune Interval** field, specify period in seconds at which PIM Join/Prune messages are transmitted on the interface.
   The range is from 0 to 18000. The default is 60.
8. From the **BSR Border** menu, select **Enable** or **Disable** to enable or disable the bootstrap router (BSR) border on the interface.
   The default is Disable.
9. In the **DR Priority** field, specify the designated router (DR) priority for the interface.
   The range is from 0 to 2147483647. The default is 1.
10. Click the **Apply** button.
    Your settings are saved.
11. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only field on the page.

Table 151.IPv6 PIM interface configuration information

Field	Description
Protocol State	The state of PIM on the interface (Operational or Non-Operational).
IPv6 Prefix/Length	The IPv6 address and prefix length of the PIM interface.
Designated Router	The IPv6 address of the designated router on the PIM interface.
Neighbor Count	The number of PIM neighbors on the PIM interface.

Display IPv6 PIM neighbors and search the PIM neighbor database

You can display the IPv6 PIM neighbor information and search the PIM neighbors database by interface or by neighbor IP address.

To display IPv6 PIM neighbors and search the PIM neighbors database:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > IPv6 PIM > PIM Neighbor**.
   The PIM Neighbor page displays.
3. To search for IGMP group entries by interface or group, select one of the following options from the **Search** menu:
   - **Interface**: Select **Interface** from the menu, specify the interface in unit/port format (for example 0/8), and click the **Go** button. If the entry exists, the entry is displayed as the first entry, followed by the remaining entries.
   - **Neighbor IP**: Select **Neighbor IP** from the menu, specify the neighbor IPv6 address, and click the **Go** button. If the entry exists, that entry with the matching neighbor IPv6 address is displayed as the first entry, followed by the remaining entries. An exact match is required.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 152.IPv6 PIM neighbor information

Field	Description
Interface	The interface on which the PIM neighbor is connected.
Neighbor IP	The IPv6 address of the PIM neighbor.
Up Time (hh:mm:ss)	The time that passed since the PIM device became a neighbor of the interface.
Expiry Time (hh:mm:ss)	The time remaining before the PIM neighbor ages out.
DR Priority	The designated router (DR) priority of the PIM neighbor.

Add an IPv6 PIM candidate rendezvous point configuration

You can add an IPv6 PIM candidate rendezvous point (RP) configuration on an interface.

To add an IPv6 PIM candidate RP configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > IPv6 PIM > Candidate RP Configuration**.
   The Candidate RP Configuration page displays.
3. From the **Interface** menu, select the interface.
4. In the **Group Address** field, specify the group IPv6 address that is transmitted in candidate-RP-advertisements.
5. In the **Prefix Length** field, enter the group IPv6 prefix length that is transmitted in candidate-RP-advertisements.
6. In the **C-RP Advertisement Interval** field, specify the period in seconds at which candidate-RP-advertisements are sent as unicast traffic to the bootstrap router (BSR).
   The range is from 1 to 16383 seconds.
7. Click the **Add** button.
   Your settings are saved. The PIM candidate RP is added.
8. To save the settings to the running configuration, click the **Save** icon.

Delete an IPv6 PIM candidate rendezvous point configuration

You can delete an IPv6 PIM candidate rendezvous point (RP) configuration that you no longer need.

To delete an IPv6 PIM candidate RP configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > IPv6 PIM > Candidate RP Configuration**.
   The Candidate RP Configuration page displays.
3. From the **Interface** menu, select the interface.
4. Select the check box for the PIM candidate RP.
5. Click the **Delete** button.
   Your settings are saved. The PIM candidate RP entry is removed.

Configure an interface as an IPv6 PIM bootstrap router candidate

You can configure an interface as an IPv6 PIM bootstrap router (BSR) candidate.

To configure an interface as an IPv6 PIM BSR candidate:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > IPv6 PIM > BSR Candidate Configuration**.
   The BSR Candidate Configuration page displays.
3. From the **Interface** menu, select an interface.
4. In the **Hash Mask Length** field, specify the hash mask length to be advertised in bootstrap messages.
   The hash mask length is used in the hash algorithm for selecting the RP for a particular group. The range is from 0 to 128. The default is 126.
5. In the **Priority** field, specify the priority for the BSR candidate.
   The range is from 0 to 255. The default is 0.
6. In the **Advertisement Interval** field, specify the period in seconds between advertisements.
   The range is from 1 to 16383. The default is 60.

7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.
9. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 153.IPv6 PIM BSR candidate configuration information

Field	Description
BSR Expiry Time (hh:mm:ss)	The time (in hours, minutes and seconds) after which the learned elected BSR expires.
IP Address	The IPv6 address of the elected BSR.
Next bootstrap Message (hh:mm:ss)	The time (in hours, minutes, and seconds) after which the next bootstrap message is due from the BSR.
Next Candidate RP Advertisement (hh:mm:ss)	The time (in hours, minutes, and seconds) after which the next candidate RP advertisement is sent.

Delete an IPv6 PIM bootstrap router candidate configuration

You can delete an IPv6 PIM bootstrap router (BSR) candidate configuration from an interface.

To delete an IPv6 PIM BSR candidate configuration from an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > IPv6 PIM > BSR Candidate Configuration**.
   The BSR Candidate Configuration page displays.
3. From the **Interface** menu, select the interface.
4. Click the **Delete** button.
   Your settings are saved. The PIM BSR candidate configuration is deleted from the interface.
5. To save the settings to the running configuration, click the **Save** icon.

Configure a static IPv6 PIM rendezvous point for a group

You can add a new or change an existing static IPv6 PIM rendezvous point (RP) for a group.

To add or change a static IPv6 PIM RP for a group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > IPv6 PIM > Static RP Configuration**.
   The Static RP Configuration page displays.
3. If you are changing an existing static RP, select the check box for the static RP.
4. In the **RP Address** field, specify the IPv6 address for the static RP.
5. In the **Group Address** field, specify the IPv6 address for the group to which the static RP applies.
6. In the **Prefix Length** field, specify the IPv6 prefix length for the group to which the static RP applies.
7. From the **Override** menu, select **Enable** or **Disable** to enable or disable the override capacity of a static RP over an RP that is learned from a BSR.
   If you select **Enable** and conflict occurs, a static RP overrides an RP that is learned from a BSR.
8. Do one of the following:
   - If you are adding a new static RP, click the **Add** button. Your settings are saved. The static RP is added.
   - If you are changing an existing static RP, click the **Apply** button. Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Delete a static IPv6 PIM rendezvous point configuration

You can delete a static IPv6 PIM rendezvous point (RP) configuration that you no longer need for a group.

To delete a static IPv6 PIM RP configuration for a group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > IPv6 PIM > Static RP Configuration**.
   The Static RP Configuration page displays.
3. Select the check box for the static RP.
4. Click the **Delete** button.
   Your settings are saved. The static RP configuration is deleted for the group.

5. To save the settings to the running configuration, click the **Save** icon.

MLD for IPv6 multicast routing

You can configure the Multicast Listener Discovery (MLD) settings and view the MLD statistics. MLD is a protocol used by IPv6 multicast routers to discover the presence of multicast listeners (nodes configured to receive IPv6 multicast packets) on its directly attached links and to discover which multicast packets are of interest to neighboring nodes. MLD is derived from IGMP; MLD version 1 (MLDv1) is equivalent to IGMPv2, and MLD version 2 (MLDv2) is equivalent to IGMPv3. MLD is a subprotocol of Internet Control Message Protocol version 6 (ICMPv6), and MLD messages are a subset of ICMPv6 messages, identified in IPv6 packets by a preceding Next Header value of 58.
The switch can snoop on both MLDv1 and MLDv2 protocol packets and bridge IPv6 multicast data based on destination IPv6 multicast MAC addresses. The switch can support MLD snooping and IGMP snooping simultaneously.

Configure the global MLP settings for the switch

You can configure the global MLDP settings for the switch, as opposed to the MLD settings for an interface.

To configure the global MLP settings for the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > MLD > Global Configuration**.
   The Global Configuration page displays.
3. Select the Admin Mode **Enable** or **Disable** radio button to enable or disable MLD on the switch.
   The default is Disable.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure an MLD routing interface

You can configure one or more MLD routing interfaces.

To configure an MLD routing interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > MLD > Routing Interface Configuration**.
   The Routing Interface Configuration page displays.
3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Admin Mode** menu, select **Enable** or **Disable** to enable or disable MLD for the interface.
   The default is Disable.
6. In the **Version** field, specify the version of MLD for the interface.
   The version can be V1 or V2. The default is V2.
7. In the **Query Interval** field, enter the frequency in seconds at which MLD host-query packets are to be transmitted on the interface.
   The value is from 1 to 3600. The default is 125.
8. In the **Query Max Response Time** field, enter the maximum query response time, in milliseconds, to be advertised in MLDv2 queries on the interface.
   The value is from 0 to 65535. The default is 10000 milliseconds.
9. In the **Startup Query Interval** field, enter the number of seconds between the transmission of startup queries on the interface.
   The value is from 1 to 300. The default is 31.

10. In the **Startup Query Count** field, enter the number of queries to be sent on startup, separated by the startup query interval.
    The value is from 1 to 20. The default is 2.
11. In the **Last Member Query Interval** field, enter the last member query interval in milliseconds.
    This value is the maximum response time to be inserted into group-specific queries sent in response to leave group messages, and is also the period between group-specific query messages. The value is from 0 to 655355. The default is 1000 milliseconds.
12. In the **Last Member Query Count** field, enter the number of queries to be sent on receiving a leave group report.
    The value is from 1 to 20. The default is 2.
13. Click the **Apply** button.
    Your settings are saved.
14. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 154.MLD Routing Interface Configuration

Field	Description
Operational Mode	The operational status of MLD on the Interface.
Robustness	The robustness setting for the interface. The default value is 2.

Display the statistics for the MLD routing interfaces

You can display the statics for the MLD routing interfaces.

To display the statistics for the MLD routing interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > MLD > Routing Interface Statistics**.
   The Routing Interface Statistics page displays.

3. Select whether to display physical interfaces, VLANs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Both physical interfaces and VLANs are displayed, or for a switch stack, both physical interfaces and VLANs on all switches are displayed.
4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 155.MLD routing interface statistics information

Field	Description
Interface	The interface on which MLD is enabled.
Querier Status	Indicates if the interface is an MLD querier or non-querier on the IPv6 subnet it is associated with.
Querier IP	The IPv6 address of the MLD querier on the IPv6 subnet to which the interface is attached.
Querier Up Time	The period in seconds since the MLD interface querier was last changed.
Querier Expiry Time	The period in seconds remaining before the other querier present timer expires. If the switch is the querier, this is zero.
Wrong Version Queries Received	The number of queries that were received on the interface with an MLD version that does not match the MLD version configured for the interface
Number of Joins Received	The number of times a group membership was added on the interface; that is, the number of times an entry for the interface was added to the cache table.
Number of Groups	The current number of entries for the interface in the cache table.

Display the MLD groups and search the MLD group database

You can display the MLD group information and search the MLD group database by interface or by group.

To display the MLD groups and search the MLD group database:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > MLD > MLD Groups**.
   The MLD Groups page displays.
3. To search for IGMP group entries by interface or group, select one of the following options from the **Search** menu:
   - **Interface**: Select **Interface** from the menu, specify the interface in unit/port format (for example 0/8), and click the **Go** button. If the entry exists, the entry is displayed as the first entry, followed by the remaining entries.
   - **Group**: Select **Group** from the menu, specify the multicast group IPv6 address, and click the **Go** button. If the entry exists, that entry with the matching group IPv6 address is displayed as the first entry, followed by the remaining entries. An exact match is required.
4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 156.MLD groups information

Field	Description
Interface	The interface on which MLD is enabled.
Group IP	The IP multicast group address.
Last Reporter	The IP address of the source of the last membership report that was received for the IP multicast group address on the interface.
Up Time	The period elapsed since this entry was created.
Expiry Time	The period remaining before this entry is aged out.
Filter Mode	The source filter mode (Include, Exclude, or NA) for the group on the interface. If the mode is NA, the field is blank.
Version 1 Host Timer	The period remaining until the switch determines that no MLD version 1 members are attached to the interface.
Group Compat Mode	The compatibility mode (MLDv1 or MLDv2) of the multicast group on the interface.
Source Hosts	The source addresses that are members of the multicast address.
Source Address (Expiry Time)	The expiration time after which each source address that is a member of the multicast group ages out.

Display or clear MLD traffic statistics

You can view the MLD traffic statistics and clear these statistics.

To view or clear IPv6 MLD traffic statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > MLD > MLD Traffic**.
   The MLD Traffic page displays.
3. To refresh the page, click the **Refresh** button.
4. To reset (clear) all MLD traffic statics, click the **Clear** button.
5. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 157.MLD traffic statistics information

Field	Description
Valid MLD Packets Received	The number of valid MLD packets received by the switch.
Valid MLD Packets Sent	The number of valid MLD packets sent by the switch.
Queries Received	The number of valid MLD queries received by the switch.
Queries Sent	The number of valid MLD queries sent by the switch.
Reports Received	The number of valid MLD reports received by the switch.
Reports Sent	The number of valid MLD reports sent by the switch.
Leaves Received	The number of valid MLD leave messages received by the switch.
Leaves Sent	The number of valid MLD leave messages sent by the switch.

Configure an MLD proxy interface

You can configure an MLD proxy interface. The following requirements apply:

- Routing, MLD, and multicast must be globally enabled on the switch.
- You must configure at least one routing interface.
- The MLD proxy interface cannot be the same interface as an MLD routing interface.

To configure an MLD proxy interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > MLD > Proxy Interface Configuration**.
   The Proxy Interface Configuration page displays.
3. From the **Interface** menu, select the interface that must function as the IGMP proxy interface.
4. From the **Admin Mode** menu, select **Enable** or **Disable** to enable or disable the IGMP proxy capability on the interface.
   The default is Disable.
5. In the **Unsolicited Report Interval** field, specify the period in seconds between repetitions of a host's initial report of membership in a group.
   The range is from 1 to 260. The default is 1.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 158.MLD proxy interface configuration information

Field	Description
IPv6 Prefix	The IPv6 address of the MLD proxy interface.
Prefix Length	The prefix length for the IPv6 address of the MLD proxy interface.
Operational Mode	The operational state of IGMP proxy interface (Disable or Enable).
Querier Address on Proxy Interface	The querier address on the proxy interface.
Number of Groups	The number of multicast group entries for the MLD proxy interface in the cache table.
Version	The version of MLD that is configured on the MLD routing interface that you changed into the MLD proxy interface. For more information, see Configure an MLD routing interface on page 508.
Version 1 Querier Timeout	The MLDv1 querier time-out period in seconds. If a host receives an MLDv1 query, the host temporarily sets its MLDv1 querier present timer to the MLDv1 querier Interval. After the host reads the MLDv1 query, the MLDv1 querier time-out indicates the period after which the host transitions back to MLDv2.
Proxy Start Frequency	The number of times the proxy was started.

Display the statistics for the MLD proxy interface

You can display the statistics for the MLD proxy interface.

To display the statistics for the MLD proxy interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > MLD > Proxy Interface Statistics**.
   The Proxy Interface Statistics page displays.

3. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 159.MLD proxy interface statistics information

Field	Description
Proxy Interface	The interface on which MLD proxy packets are received.
Version	The version of MLD proxy packets received.
Queries Received	The number of MLD proxy queries received.
Reports Received	The number of MLD proxy reports received.
Reports Sent	The number of MLD proxy reports sent.
Leaves Received	The number of MLD proxy leave messages received.
Leaves Sent	The number of MLD proxy leave messages sent.

Display the MLD proxy membership and search the MLD proxy membership database

You can display the MLD proxy membership information and search the MLD proxy membership database by group.

To display the MLD proxy membership information and search the MLD proxy membership database:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > MLD > Proxy Membership**.
   The Proxy Membership page displays.
3. To search for MLD proxy membership by group, specify the multicast group IPv6 address, and click the **Go** button.
   If the entry exists, that entry with the matching group IPv6 address is displayed as the first entry, followed by the remaining entries. An exact match is required.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 160.MLD proxy membership information

Field	Description
Proxy Interface	The interface on which the MLD proxy is enabled.
Group IP	The IPv6 multicast group address.
Source Hosts	The source addresses that are members of the multicast group.
Last Reporter	The source IPv6 address that sent the last membership report for the multicast group on the MLD proxy interface.
Up Time	The time elapsed since this entry was created.
Expiry Time	The expiration interval for each source address that is a member of the multicast group. This is the period after which the source entry ages out.
State	The state of the host entry: ○ **Non-member state**: The host does not belong to the group on the interface. ○ **Delaying member state**: The host belongs to the group on the interface and its report timer is active. (The report timer is used to send reports.) ○ **Idle member state**: The host belongs to the group on the interface but its report timer is inactive.
Filter Mode	The group filter mode (Include, Exclude, or None) for the group on the MLD proxy interface.
Number of Sources	The number of source hosts in the multicast group.

Static multicast routes for IPv6 addresses

If a multicast routing protocol such as PIM-SM or PIM-DM is enabled and configured, multicast routes are built dynamically. However, you can also configure one or more static IPv6 multicast routes. A static route can facilitate the processing of multicast traffic when a neighboring reverse-path forwarding (RPF) router is located between the switch and the multicast source.

Configure static multicast routes for IPv6 addresses

You can add a new or change an existing static IPv6 multicast route.

To add or change a static IPv6 multicast route:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > Static Routes Configuration**.
   The Static Routes Configuration page displays.
3. If you are changing an existing static multicast route, select the check box for the route.
   For an existing static route, you can change the metric only.
4. In the **Source IP** field, specify the IPv6 address of the multicast packet source.
5. In the **Prefix Length** field, specify the IPv6 prefix length of the multicast packet source.
6. In **RPF Neighbor** field, specify the IPv6 address of the neighbor RPF router that is located between the switch and the multicast source.
7. In the **Metric** field, specify the link state cost of the path from the multicast source to the switch.
   The range is from 0 to 255.
8. From the **RPF Interface** menu, select the interface that is connected to the RPF neighbor.
9. Do one of the following:
   - If you are adding a new static multicast route, click the **Add** button. Your settings are saved. The static multicast route is added.
   - If you are changing an existing static multicast route, click the **Apply** button. Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

Delete a static multicast route for an IPv6 address

You can delete a static IPv6 multicast route that you no longer need.

To delete a static multicast route:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Routing > IPv6 Multicast > Static Routes Configuration**.
   The Static Routes Configuration page displays.
3. Select the check box for the static route.
4. Click the **Delete** button.
   Your settings are saved. The static route is deleted.

5. To save the settings to the running configuration, click the **Save** icon.

Configure Quality of Service

This chapter covers the following topics:

- Quality of Service concepts
- Class of Service
- Differentiated Services

Quality of Service concepts

In a switch, each physical port consists of one or more queues for transmitting packets on the attached network. Multiple queues per port are often provided to give preference to certain packets over others based on user-defined criteria. When a packet is queued for transmission in a port, the rate at which it is serviced depends on how the queue is configured and possibly the amount of traffic present in the other queues of the port. If a delay is necessary, packets are held in the queue until the scheduler authorizes the queue for transmission. As queues become full, packets can no longer be held for transmission and are dropped by the switch.
Quality of Service (QoS) is a means of providing consistent, predictable data delivery by distinguishing packets with strict timing requirements from those that are more tolerant of delay. Packets with strict timing requirements are given special treatment in a QoS-capable network. With this in mind, all elements of the network must be QoS capable. The presence of at least one node that is not QoS capable creates a deficiency in the network path, and the performance of the entire packet flow is compromised.

Class of Service

The Class of Service (CoS) queueing feature lets you directly configure certain aspects of switch queueing. This provides the desired QoS behavior for different types of network traffic when the complexities of DiffServ are not required. The priority of a packet arriving at an interface can be used to steer the packet to the appropriate outbound CoS queue through a mapping table. CoS queue characteristics that affect queue mapping, such as minimum guaranteed bandwidth or transmission rate shaping, are user-configurable at the queue (or port) level.
Eight queues per port are supported.

CoS configuration concepts

You can set the Class of Service trust mode for an interface. Each port on the switch can be configured to trust one of the packet fields (802.1p or IP DSCP), or to not trust any packet's priority designation (untrusted mode). If the port is set to a trusted mode, it uses a mapping table appropriate for the trusted field being used. This mapping table indicates the CoS queue to which the packet must be forwarded on the appropriate egress port. Of course, the trusted field must

exist in the packet for the mapping table to be of any use. If this is not the case, default actions are performed. These actions involve directing the packet to a specific CoS level configured for the ingress port as a whole, based on the existing port default priority as mapped to a traffic class by the current 802.1p mapping table.
Alternatively, when a port is configured as untrusted, it does not trust any incoming packet priority designation and uses the port default priority value instead. All packets arriving at the ingress of an untrusted port are directed to a specific CoS queue on the appropriate egress ports, in accordance with the configured default priority of the ingress port. This process is also used for cases where a trusted port mapping cannot be honored, such as when a non-IP packet arrives at a port configured to trust the IP DSCP value.

Configure the CoS trust mode settings globally or for a specific interface

A global CoS trust mode configuration setting is applied to all interfaces on the switch.
A CoS trust mode configuration setting for an interface is applied to the selected interface only. To configure the CoS trust mode for more than one interface, repeat the task for each selected interface.

To configure the CoS trust mode settings on all interfaces or on a specific interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > CoS > Basic > CoS Configuration**.
   The CoS Configuration page displays.

3. Either configure the same CoS trust mode settings for all CoS-configurable interfaces or configure CoS settings per interface.
   By default, the Global radio button is selected.
   - To configure the same CoS trust mode settings for all CoS configurable interfaces, do the following:
     a. Select the **Global** radio button.
     b. From the **Global Trust Mode** menu, select one of the following trust mode options for ingress traffic on the switch:
        - **untrusted**: Do not trust any CoS packet marking at ingress.
        - **trust dot1p**: The eight priority tags that are specified in IEEE 802.1p are p0 to p7. The QoS setting lets you map each of the eight priority levels to one of eight internal hardware priority queues.
        - **trust IP-DSCP**: The six most significant bits of the DiffServ field are called the Differentiated Services Code Point (DSCP) bits. The default mode is IP-DSCP.
   - To configure CoS settings for an interface, do the following:
     a. Select the **Interface** radio button.
     b. From the menu, select the interface. The configured settings apply to the selected interface only.
     c. From the **Interface Trust Mode** menu, select one of the following trust mode options:
        - **untrusted**: Do not trust any CoS packet marking at ingress.
        - **trust dot1p**: The eight priority tags that are specified in IEEE 802.1p are p0 to p7. The QoS setting lets you map each of the eight priority levels to one of seven internal hardware priority queues.
        - **trust IP-DSCP**: The six most significant bits of the DiffServ field are called the Differentiated Services Code Point (DSCP) bits. The default mode is IP-DSCP.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Map 802.1p priorities to queues

You can view or change which internal traffic classes are mapped to the 802.1p priority class values in Ethernet frames that the device receives. The priority-to-traffic class mappings can be applied globally (to all interfaces) or to each individual interface. The mapping allows the switch to group various traffic types (for example, data or voice) based on their latency requirements and give preference to time-sensitive traffic.

To map 802.1p priorities to queues for all interfaces or individual interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > CoS > Advanced > 802.1p Queue Mapping**.
   The 802.1p Queue Mapping page displays.
3. From the **Interface** menu, either select **All** to apply the mapping to all interfaces or select an individual interface.
4. In the 802.1p to Queue Mapping table, map each of the eight 802.1p priorities to a queue (internal traffic class) from **0** to **6**.
   The 802.1p Priority row contains traffic class selectors for each of the eight 802.1p priorities to be mapped. The priority goes from low (0) to high (7). For example, traffic with a priority of 0 is for most data traffic and is sent using best effort. Traffic with a higher priority, such as 7, might be time-sensitive traffic, such as voice or video.

   Table 161.Default values for 802.1p to queue mapping

802.1p priority	0	1	2	3	4	5	6	7
Default queue	1	0	0	1	2	2	3	3

   The values in the menu under each priority represent the traffic class. The traffic class is the hardware queue for a port. Higher traffic class values indicate a higher queue position. Before traffic in a lower queue is sent, it must wait for traffic in higher queues to be sent.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Map DSCP values to queues

You can map an internal traffic class to a DSCP value.

To map DSCP values to queues:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > CoS > Advanced > IP DSCP Queue Mapping**.
   The IP DSCP Queue Mapping page displays.
   The IP DSCP field displays an IP DSCP value from 0 to 63.
3. For each DSCP value, specify which internal traffic class to map the corresponding queue.

4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the CoS interface settings for an interface

You can configure the trust mode for one or more interfaces and apply an interface shaping rate to all interfaces or to a specific interface.

To configure CoS settings for an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > CoS > Advanced > CoS Interface Configuration**.
   The CoS Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

5. From the **Interface Trust Mode** menu, select one of the following trust mode options for ingress traffic on the selected interfaces:
   - **Untrusted**: Do not trust any CoS packet marking at ingress.
   - **802.1p**: The eight priority tags that are specified in IEEE 802.1p are p0 to p7. The QoS setting lets you map each of the eight priority levels to one of seven internal hardware priority queues. The default value is 802.1p.
   - **IP-DSCP**: The six most significant bits of the DiffServ field are called the Differentiated Services Code Point (DSCP) bits.
6. In the **Interface Shaping Rate** field, specify the maximum outbound transmission rate bandwidth in kbps.
   This setting is used to shape the outbound transmission rate in increments of 1 percent in a range of 0–100. This value is controlled independently of any per-queue maximum bandwidth configuration. It is effectively a second-level shaping mechanism. The default value is 0. The value 0 means that the maximum is unlimited.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Configure CoS queue settings for an interface

You can define what a particular queue does by configuring switch egress queues. You can control how much bandwidth is used by the queue, the queue depth during times of congestion, and the scheduling of packet transmission from all queues on a port. Each port contains its own CoS queue-related configuration.
For information about configuring CoS queue settings globally, see Configure the CoS trust mode settings globally or for a specific interface on page 521.

To configure the CoS queue settings for an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > CoS > Advanced > Interface Queue Configuration**.
   The Interface Queue Configuration page displays.

3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Queue ID** menu, select the queue.
   You can select a queue from **0** to **6**. (For CoS, queue 7 is reserved for stacking.)
6. In the **Minimum Bandwidth** field, specify the minimum guaranteed bandwidth allotted to the queue on the interface.
   Setting this value higher than its corresponding maximum bandwidth automatically increases the maximum to the same value. The default value is 0. The range is from 0 to 100 in increments of 1. The value 0 means that no guaranteed minimum bandwidth is applied. The sum of the individual minimum bandwidth values for all queues on the selected interface cannot exceed the maximum (100).
   The setting applies only to the queue that you select from the **Queue ID** menu.
7. From the **Scheduler Type** menu, select the type of scheduling that is used for the queue on the interface:
   - **Strict**: The interface services traffic with the highest priority on a queue first.
   - **Weighted**: The interface uses weighted round robin to associate a weight to each queue. This is the default setting.

   The setting applies only to the queue that you select from the **Queue ID** menu.

8. From the **Queue Management Type** menu, select the queue depth management method that is used for the queue on the interface:
   - **TailDrop**: The Taildrop method causes packets to be dropped if the queue is full. When space becomes available in the queue, packets are once again accepted. This is the default method.
   - **WRED**: The Weighted Random Early Detection (WRED) method allows you to refine queue management so that packets are dropped selectively (see Configure the CoS WRED precedence settings for dropping packets on page 527).

   The setting applies only to the queue that you select from the **Queue ID** menu.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

Configure the CoS WRED precedence settings for dropping packets

If you use the Weighted Random Early Detection (WRED) method as the queue depth management method (see Configure CoS queue settings for an interface on page 525), you can configure the precedence settings for dropping packets if a queue is full. These settings apply to a selected interface and a selected queue. You can configure different setting for each queue and for each precedence level on one interface.

To configure CoS WRED precedence settings for dropping packets:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > CoS> Advanced > CoS Queue Drop Precedence Configuration**.
   The CoS Queue Drop Precedence Configuration page displays.
3. From the **Interface** menu, select an interface.
4. From the **Queue ID** menu, select a queue from **0** to **6**. (For CoS, queue 7 is reserved for stacking.)

5. From the **Drop Precedence Level** menu, select the drop precedence levels from **1** to **4**.
   Packets to which level 1 applies are the most likely to be dropped. Packets to which level 4 applies are the least likely to be dropped.
   The following table shows the default values for each drop precedence level. You can modify these settings.

   Table 162.Default values for drop precedence levels

Drop precedence level	WRED Minimum Threshold	WRED Maximum Threshold	WRED Drop Probability Scale
1	40	100	10
2	30	90	10
3	20	80	10
4	99	100	10

6. In the **WRED Minimum Threshold** field, specify the WRED minimum queue threshold below which no packets are dropped for the selected drop precedence level.
   The range is from 0 to 100.
7. In the **WRED Maximum Threshold** field, specify the weighted RED maximum queue threshold above which all packets are dropped for the current drop precedence level.
   The range is from 0 to 100.
8. Use **WRED Drop Probability Scale** field, specify the packet drop probability for the current drop precedence level.
   The range is from 0 to 100.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 163.CoS Interface Queue Drop Precedence Status

Field	Description
Interface	The CoS interface.
Queue ID	The queue ID.
Drop Precedence Level	The drop precedence level.
WRED Minimum Threshold	The WRED minimum queue threshold value.
WRED Maximum Threshold	The WRED maximum queue threshold value.
WRED Drop Probability Scale	The WRED packet drop probability value.

Differentiated Services

The QoS feature contains Differentiated Services (DiffServ) support that allows traffic to be classified into streams and given certain QoS treatment in accordance with defined per-hop behaviors.

Standard IP-based networks provide best-effort data delivery service. Best-effort service implies that the network delivers the data in a timely fashion, although there is no guarantee. If congestion occurs, packets might be delayed, sent sporadically, or dropped. For typical Internet applications, such as email and file transfers, a slight degradation in service is acceptable and in many cases unnoticeable. However, any degradation of service can negatively affect applications with strict timing requirements, such as voice and multimedia.

Defining DiffServ

To use DiffServ for QoS, you must first define the following categories and their criteria:

1. **Class**: Create classes and define class criteria.
2. **Policy**: Create policies, associate classes with policies, and define policy statements.
3. **Service**: Add a policy to an inbound interface.

Packets are filtered, classified, and processed based on defined criteria. The filtering and classification criteria are defined by a class. The processing is defined by a policy's attributes. Policy attributes can be defined on a per-class instance basis, and it is these attributes that are applied when a match occurs. A policy can contain multiples classes. When the policy is active, the actions taken depend on which class matches the packet.

You begin the DiffServ configuration by defining one or more match criteria for a class. Then, you set up a policy and add one or more classes to the policy. Finally, you add a service, which means that you add a policy to an inbound interface.

After the DiffServ configuration is complete, the switch begins packet processing by testing the

match criteria for a packet:

- **All class type**: The All class type option specifies that each match criteria within a class must be true for a packet to match that class.
- **Any class type**: The Any class type option specifies that at least one match criteria must be true for a packet to match that class.

Classes are tested in the order in which you added them to the policy. A policy is applied to a packet when a class match within that policy is found.

DiffServ wizard overview

The DiffServ wizard enables DiffServ on the switch by creating a traffic class, adding the traffic class to a policy, and then adding the policy to the ports and LAGs that you select. The DiffServ wizard does the following:

- Creates a DiffServ class and defines match criteria used as a filter to determine if incoming traffic meets the requirements to be a member of the class.
- Sets the DiffServ class match criteria based on traffic type selection as follows:
  - **VOIP**: Sets the match criteria for VoIP traffic to the UDP protocol.
  - **HTTP**: Sets the match criteria for HTTP traffic to the HTTP destination port.
  - **FTP**: Sets match criteria for FTP traffic to the FTP destination port.
  - **Telnet**: Sets the match criteria for Telnet traffic to the Telnet destination port.
  - **Every**: Sets the match criteria for all traffic.
- Sets a committed rate for the selected traffic type.
- Creates a Diffserv policy and adds it to the DiffServ class.
- Optionally, enables policing and sets the Diffserv policy as a simple policy. When policing is enabled, traffic that conforms to the class match criteria is processed according to the outbound priority selection, which handles conforming traffic as follows:
  - **High**: Maps the traffic priority to DiffServ class Expedited Forwarding (EF).
  - **Medium**: Maps the traffic priority to DiffServ class Assured Forwarding (AF31).
  - **Low**: Maps the traffic priority to DiffServ class Best Effort (BE).

  When policing is disabled, the outbound priority selection (high, medium, or low) affects *all* traffic.
- Optionally, adds selected ports and LAGs to the policy.

Use the DiffServ wizard to create a traffic class and policy for one or more interfaces

To use the DiffServ wizard to create a traffic class and policy for one or more interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **QoS > DiffServ > DiffServ Wizard**.
   The DiffServ Wizard page displays.
3. From the **Traffic Type** menu, define the DiffServ class by selecting **VOIP**, **HTTP**, **FTP**, **Telnet**, or **Every**.
4. In the **Committed Rate (Kbps)** field, specify the rate in Kbps that applies to the traffic in the class.
   The range is from 1 to the maximum supported bandwidth on the interface. The default is 1 Kbps.
   The committed rate is applied if policing is enabled (which it is by default)
5. To disable policing, clear the **Policing** check box.
   By default, policing is enabled and the outbound priority (high, medium, or low) is set as a simple policy. For more information, see the following step.
   When policing is disabled, the outbound priority (high, medium, or low) affects *all* traffic.
6. From the **Outbound Priority** menu, select the priority for outbound traffic:
   - **High**: Maps the traffic priority to DiffServ class Expedited Forwarding (EF).
   - **Medium**: Maps the traffic priority to DiffServ class Assured Forwarding (AF31).
   - **Low**: Maps the traffic priority to DiffServ class Best Effort (BE).

   Traffic that conforms to the class match criteria is processed according to the outbound priority selection.
7. In the Ports table (or if a stack is configured, in one or more of the Ports tables), click each port to which you want to add the DiffServ policy.
   The ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed, and you can select ports on different stacked switches.
8. In the LAG table, click each LAG to which you want to add the DiffServ policy.
   All LAGs are displayed, whether or not a stack is configured.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

Configure the DiffServ mode and display the entries in the DiffServ private MIB tables

You can enable or disable DiffServ and display the current and maximum number of rows in each of the main DiffServ private MIB tables.

To configure the DiffServ mode and display the entries in the DiffServ private MIB tables:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Basic > DiffServ Configuration**.
   The DiffServ Configuration page displays.
3. Select the administrative mode for DiffServ:
   - **Enable**: Differentiated services are active. This is the default setting.
   - **Disable**: Differentiated services are inactive. An existing DiffServ configuration is retained and can be changed but is not active.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 164.DiffServ Status information

Field	Description
Class table	The number of configured DiffServ classes out of the total allowed on the switch.
Class Rule table	The number of configured class rules out of the total allowed on the switch.
Policy table	The number of configured policies out of the total allowed on the switch.
Policy Instance table	The number of configured policy class instances out of the total allowed on the switch.
Policy Attributes table	The number of configured policy attributes (attached to the policy class instances) out of the total allowed on the switch.
Service table	The number of configured services (attached to the policies on specified interfaces) out of the total allowed on the switch.

Configure a DiffServ class

You can add a new DiffServ class name or rename or delete an existing class. You can also define the criteria to associate with a DiffServ class. As packets are received, these DiffServ classes are used to prioritize packets. You can use multiple match criteria in a class. The logic is a Boolean logical-AND for this criteria.

Add and configure a DiffServ class

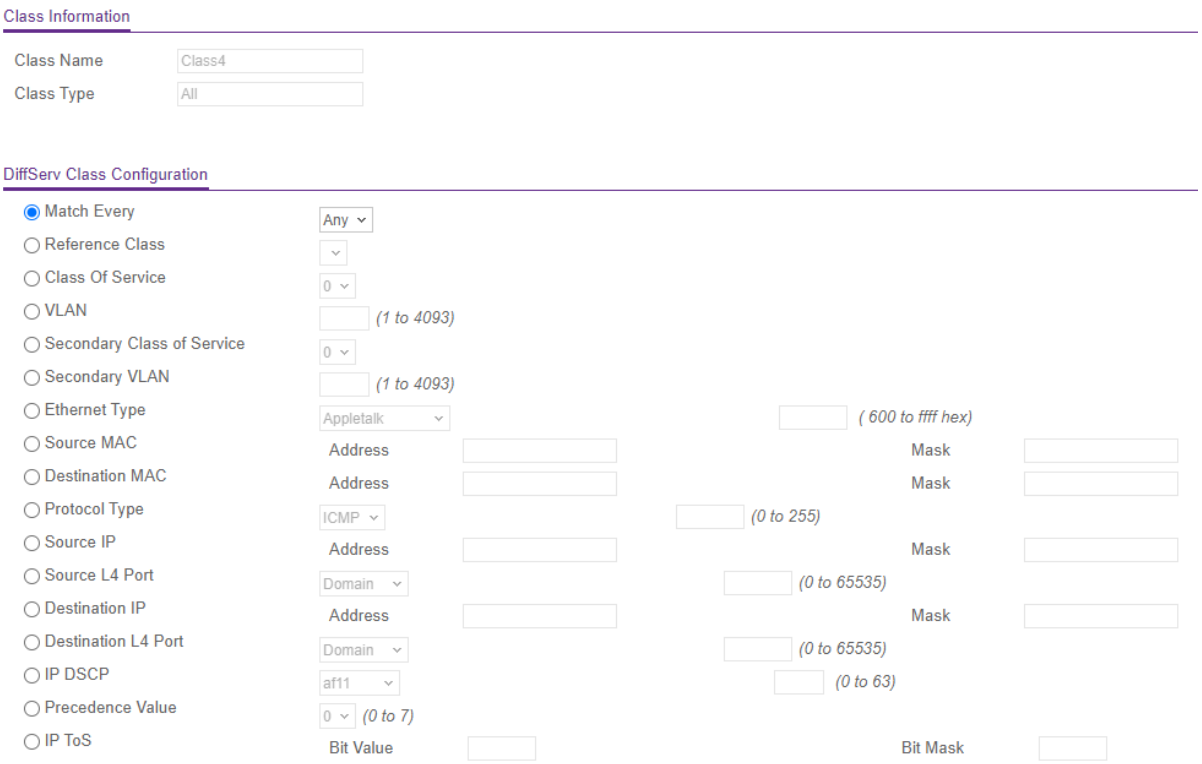
You can add a DiffServ class and configure the criteria that must be associated with it.

To add and configure a DiffServ class:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > Class Configuration**.
   The Class Configuration page displays.
3. In the **Class Name** field, enter a class name.
   The **Class Name** field also lists all existing DiffServ class names, from which you can select one for modification or deletion. The class name can be 1 to 31 alphanumeric characters in length.
4. From the **Class Type** menu, select one of the following options:
   - **All**: All match criteria that you define for the class must be satisfied for a packet match. All signifies the logical AND of all the match criteria. For example, if the class includes one criterion for an IP address and another criterion for a MAC address, the traffic must match both criteria.
   - **Any**: Any match criterion that you define for the class must be satisfied for a packet match. For example, if the class includes one criterion for an IP address and another criterion for a MAC address, the traffic must match either one of the criteria but does not need to match both criteria.

   You can select the class type only when you are creating a new class. After the class is created, you cannot modify the type of class.
5. Click the **Add** button.
   The new class is added.

6. After creating the class, click the class name.
   The class name is a hyperlink to the page on which you can define the class configuration (see the following figure).



7. Define the criteria that must be associated with the DiffServ class by selecting *one* of the following radio buttons:
   - **Match Every**: Select this radio button to add a match condition that considers all packets to belong to the class. The only selection from the **Match Every** menu is **Any**.
   - **Reference Class**: Select this radio button to reference another class for criteria. The match criteria defined in the reference class function as match criteria in addition to the match criteria that you define for the selected class. After you select the radio button, the classes that can be referenced are displayed. Select the class to reference. A class can reference only one other class of the same type.
   - **Class of Service**: Select this radio button to require the Class of Service (CoS) value in an Ethernet frame header to match the specified CoS value. This option lists all the values for the Class of Service match criterion in the range 0 to 7 from which you can select one.
   - **VLAN**: Select this radio button to require a packet's VLAN ID to match a VLAN ID. The VLAN value is in the range from 1 to 4093.
   - **Secondary Class of Service**: Select this radio button to require the secondary Class of Service (CoS) value in an Ethernet frame header to match the specified secondary CoS value.
   - **Secondary VLAN**: Select this radio button to require a packet's VLAN ID to match a secondary VLAN ID. The VLAN value is in the range from 1 to 4093.
   - **Ethernet Type**: Select this radio button to require the EtherType value in the Ethernet frame header to match the specified EtherType value. After you select the radio button, select the EtherType keyword from the menu of common protocols that are mapped to their Ethertype value. You can also select **User Value** from the menu and enter a value in the hexadecimal range from 600 to ffff.
   - **Source MAC**: Select this radio button to require a packet's source MAC address to match the specified MAC address. After you select this radio button, use the following fields to configure the source MAC address match criteria:
     - **Address**: The source MAC address to match. The source MAC address is specified as six two-digit hexadecimal numbers separated by colons.
     - **Mask**: The MAC mask, which specifies the bits in the source MAC address to compare against the Ethernet frame. Use Fs and zeros to configure the MAC mask. An F means that the bit is checked, and a zero in a bit position means that the data is not significant. For example, if the MAC address is aa:bb:cc:dd:ee:ff, and the mask is ff:ff:00:00:00:00, all MAC addresses with aa:bb:xx:xx:xx:xx result in a match (where x is any hexadecimal number). Note that this is not a wildcard mask, which ACLs use.
   - **Destination MAC**: Select this radio button to require a packet's destination MAC address to match the specified MAC address. After you select the radio button, use the following fields to configure the destination MAC address match criteria:

        - **Address**: The destination MAC address to match. The destination MAC address is specified as six two-digit hexadecimal numbers separated by colons.
        - **Mask**: The MAC mask, which specifies the bits in the destination MAC address to compare against an Ethernet frame. Use Fs and zeros to configure the MAC mask. An F means that the bit is checked, and a zero in a bit position means that the data is not significant. For example, if the MAC address is aa:bb:cc:dd:ee:ff, and the mask is ff:ff:00:00:00:00, all MAC addresses with aa:bb:xx:xx:xx:xx result in a match (where x is any hexadecimal number). Note that this is not a wildcard mask, which ACLs use.
    - **Protocol Type**: Select this radio button to require a packet's Layer 4 protocol to match the specified protocol, which you must select from the menu. You can also select **Other** from the menu and enter a protocol number from 0 to 255.
    - **Source IP**: Select this radio button to require a packet's source IP address to match the specified IP address. After you select the radio button, use the following fields to configure the source IP address match criteria:
        - **Address**: The source IP address format to match in dotted-decimal.
        - **Mask**: The bit mask in IP dotted-decimal format indicating which parts of the source IP address to use for matching against packet content.
    - **Source L4 Port**: Select this radio button to require a packet's TCP/UDP source port to match the specified protocol, which you must select from the menu. You can also select **Other** from the menu and enter a port number from 0 to 65535.
    - **Destination IP**: Select this radio button to require a packet's destination IP address to match the specified IP address. After you select the radio button, use the following fields to configure the destination IP address match criteria:
        - **Address**: The destination IP address format to match in dotted-decimal.
        - **Mask**: The bit mask in IP dotted-decimal format indicating which parts of the destination IP address to use for matching against packet content.
    - **Destination L4 Port**: Select this radio button to require a packet's TCP/UDP destination port to match the specified protocol. You can also select **Other** from the menu and enter a port number from 0 to 65535.
    - **IP DSCP**: Select this radio button to require the packet's IP DiffServ Code Point (DSCP) value to match the specified IP DSCP keyword code, which you must select from the menu. You can also select **Other** from the menu and enter an IP DSCP value from 0 to 63. The DSCP value is defined as the high-order 6 bits of the Service Type octet in the IP header.
    - **Precedence Value**: Select this radio button to require the packet's IP precedence value to match the specified number from 0 to 7, which you must select from the menu. The IP Precedence field in a packet is defined as the high-order 3 bits of the Service Type octet in the IP header.
    - **IP ToS**: Select this radio button to require the packet's Type of Service (ToS) bits in the IP

header to match the specified value. The IP ToS field in a packet is defined as all 8 bits of the service type octet in the IP header. After you select the radio button, use the following fields to configure the ToS match criteria:

- **Bit Value**: Enter a two-digit hexadecimal number octet value in the range from 00 to ff to match the bits in a packet's ToS field.
- **Bit Mask**: Specify the bit positions that are used for comparison against the IP ToS field in a packet.

8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields in the Class Summary section.

Table 165.DiffServ Class Configuration, Class Summary information

Field	Description
Match Criteria	The configured match criteria for the specified class.
Values	The values of the configured match criteria.

Rename an existing DiffServ class

You can change the name of an existing DiffServ class.

To rename an existing DiffServ class:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > Class Configuration**.
   The Class Configuration page displays.
3. Select the check box next to the class name.
4. In the **Class Name** field, specify the new name.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Change the criteria for an existing DiffServ class

You can change the criteria for an existing DiffServ class.

To change the criteria for an existing DiffServ class:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > Class Configuration**.
   The Class Configuration page displays.
3. Click the class name, which is a hyperlink.
   The page on which you can change the class configuration displays.
4. Change the class configuration as needed.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete a DiffServ class

You can delete an existing DiffServ class that you no longer need.

To delete a DiffServ class:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > Class Configuration**.
   The Class Configuration page displays.
3. Select the check box next to the class name.
4. Click the **Delete** button.
   The class is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure an IPv6 DiffServ class

The switch supports QoS ACL and DiffServ functionality for IPv6 by providing support for IPv6 packet classification. An IPv6 ACL serves the same purpose as an IPv4 ACL.
An Ethernet IPv6 packet is distinguished from an IPv4 packet by its unique Ethertype value, so all IPv6 classifiers include the Ethertype field, even though you cannot configure its value for an IPv6

class on the switch.
The destination and source IPv6 addresses use a prefix length value instead of an individual mask to qualify them as a subnet addresses or a host addresses. Packets that match an IPv6 classifier can be marked with the IP DSCP field in the traffic class octet.
You can also assign an IPv6 ACL with a DiffServ assignment to LAG interfaces.
You can add a new DiffServ class name or rename or delete an existing class. As packets are received, these DiffServ classes are used to prioritize packets. You can use multiple match criteria in a class. The logic is a Boolean logical-AND for this criteria.

Add and configure an IPv6 DiffServ class

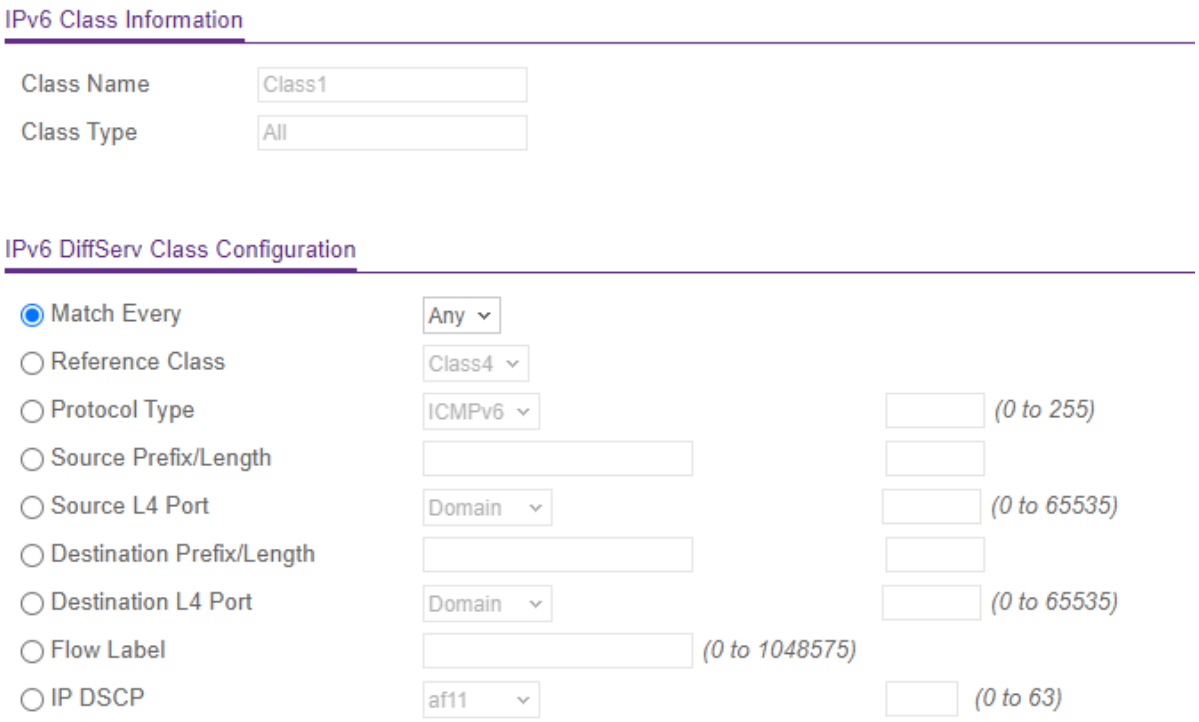
You can add an IPv6 DiffServ class and configure the criteria that must be associated with it.

To configure an IPv6 DiffServ class:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > IPv6 Class Configuration**.
   The Class Configuration page displays.
3. In the **Class Name** field, enter a class name.
   The **Class Name** field also lists all existing DiffServ class names, from which you can select one for modification or deletion. The class name can be 1 to 31 alphanumeric characters in length.
4. From the **Class Type** menu, select one of the following options:
   - **All**: All match criteria that you define for the class must be satisfied for a packet match. All signifies the logical AND of all the match criteria. For example, if the class includes one criterion for an IP address and another criterion for a MAC address, the traffic must match both criteria.
   - **Any**: Any match criterion that you define for the class must be satisfied for a packet match. For example, if the class includes one criterion for an IP address and another criterion for a MAC address, the traffic must match either one of the criteria but does not need to match both criteria.

   You can select the class type only when you are creating a new class. After the class is created, you cannot modify the type of class.
5. Click the **Add** button.
   The new class is added.

6. After creating the class, click the class name.
   The class name is a hyperlink to the page on which you can define the class configuration (see the following figure).



7. Define the criteria to associate with a DiffServ class:
   - **Match Every**: Select this radio button to add a match condition that considers all packets to belong to the class. The only selection from the **Match Every** menu is **Any**.
   - **Reference Class**: Select this radio button to reference another class for criteria. The match criteria defined in the reference class function as match criteria in addition to the match criteria that you define for the selected class. After selecting this option, the classes that can be referenced are displayed. Select the class to reference. A class can reference one other class of the same type.
   - **Protocol Type**: Select this radio button to require a packet's Layer 4 protocol to match the specified protocol, which you must select from the menu. The menu includes **Other** as a selection, which lets you enter a protocol number from 0 to 255.
   - **Source Prefix/Length**: Select this radio button to require a packet's source prefix and prefix length to match the specified source IPv6 prefix and prefix length. Prefix must always be specified with the prefix length. The prefix can be in the hexadecimal range from 0 to FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF and the prefix length can be in the range from 0 to 128.
   - **Source L4 Port**: Select this radio button to require a packet's TCP/UDP source port to match the specified protocol, which you must select from the menu. The range is 0 to 65535. The menu includes **Other** as an option for unnamed ports.
   - **Destination Prefix/Length**: Select this radio button to require a packet's destination prefix and prefix length to match the specified source IPv6 prefix and prefix length. Prefix must always be specified with the prefix length. The prefix can be in the hexadecimal range from 0 to FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF and the prefix length can be in the range from 0 to 128.
   - **Destination L4 Port**: Select this radio button to require a packet's TCP/UDP destination port to match the specified protocol, which you must select from the menu. The range is 0 to 65535. The menu includes **Other** as an option for unnamed ports.
   - **Flow Label**: Select this radio button to require a packet's flow label to match the specified flow label. The flow label is a 20-bit number that is unique to an IPv6 packet and that is used by end stations to signify QoS handling in routers. The flow label can be specified in the range from 0 to 1048575.
   - **IP DSCP**: Select this radio button to require the packet's IP DiffServ Code Point (DSCP) value to match the specified IP DSCP keyword code, which you must select from the menu. The menu includes **Other** as a selection, which lets you enter an IP DSCP value from 0 to 63. The DSCP value is defined as the high-order 6 bits of the Service Type octet in the IP header.

8. Click the **Apply** button.
   Your settings are saved.

9. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only information displayed in the Class Summary section.

Table 166.IPv6 DiffServ class configuration class summary

Field	Description
Match Criteria	The configured match criteria for the specified class.
Values	The values of the configured match criteria.

Rename an existing IPv6 DiffServ class

You can change the name for an existing IPv6 DiffServ class.

To rename an existing IPv6 DiffServ class:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > IPv6 Class Configuration**.
   The Class Configuration page displays.
3. Select the check box next to the class name.
4. In the **Class Name** field, specify the new name.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Change the criteria for an existing IPv6 DiffServ class

You can change the criteria for an existing IPv6 DiffServ class.

To change the criteria for an existing IPv6 DiffServ class:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > IPv6 Class Configuration**.
   The Class Configuration page displays.
3. Click the class name, which is a hyperlink.
   The page on which you can change the class configuration displays.

4. Change the class configuration as needed.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete an IPv6 DiffServ class

You can delete an IPv6 DiffServ class that you no longer need.

To delete an IPv6 DiffServ class:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > IPv6 Class Configuration**.
   The Class Configuration page displays.
3. Select the check box next to the class name.
4. Click the **Delete** button.
   The class is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure a DiffServ policy

You can associate a collection of classes with one or more policies.

Add and configure a DiffServ policy

You can add a policy, configure the type of the policy, add a member class to the policy, and then configure the policy attributes.

To add and configure a DiffServ policy:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > Policy Configuration**.
   The Policy Configuration page displays.
3. In the **Policy Name** field, enter a policy name.
   The policy class name can be 1 to 31 alphanumeric characters in length.

4. From the **Policy Type** menu, select the traffic flow direction to which the policy must be applied:
   - **In**: The policy is applied to inbound traffic.
   - **Out**: The policy is applied to outbound traffic.
5. From the **Member Class** menu, select an existing class that you want to associate with the new policy.
6. Click the **Add** button.
   The new policy is added.
7. After creating the policy, click the policy name.
   The policy name is a hyperlink.
   The page that displays lets you define the policy attributes.

8. Configure the policy attributes by selecting *one* of the following radio buttons:
   - **Assign Queue**: Select this radio button to specify that traffic must be assigned to a queue, which you must select from the menu. The queue is expressed as a value in the range from 0 to 7.
   - **Drop**: Select this radio button to require each packet to be dropped.
   - **Mark VLAN CoS**: Select this radio button to specify the VLAN priority, which you must select from the menu. The VLAN priority is expressed as a value in the range from 0 to 7.
   - **Mark IP Precedence**: Select this radio button to require packets to be marked with an IP precedence value before being forwarded. You must select an IP precedence value from 0 to 7 from the menu.
   - **Mirror**: Select this radio button to require packets to be mirrored to an interface or LAG, one of which you must select from the menu.
   - **Redirect**: Select this radio button to require packets to be redirected to an interface or LAG, one of which you must select from the menu.
   - **Mark IP DSCP**: Select this radio button to require packet to be marked with an IP DSCP keyword code, which you must select from the menu. The DSCP value is defined as the high-order 6 bits of the Service Type octet in the IP header.
   - **Simple Policy**: Select this radio button to define that matching packets are treated by a simple traffic policy, which is color blind and for which color classes do not apply. A simple policy supports a single data rate and results in an action for packets that conforms to the policy and an action for packets that violate the policy. You can configure the action for packets that conform to the policy, but the action for packets that violate the policy is always Drop, that is, these packets are always dropped.
     a. Specify the traffic rate and bust size for packets that conform to the simple policy.
        - **Committed Rate**. From the menu, select **Percent** or **Kbps**, and then specify the committed rate that is applied to conforming packets by entering a value in the range from 1 to 100 percent or from 1 to 4294967295 kilobits-per-second (Kbps). The committed rate is used to limit the arrival rate of conforming traffic.
        - **Committed Burst Size**. Specify the burst size that is applied to conforming packets by entering a value in the range from 1 to 128 KBytes (KB). The committed burst size is used to limit the amount of conforming traffic.
     b. **Confirm Action**: Specify a policy action for packets that conform to the simple policy.
        - **Send**: Packets are forwarded unmodified. This is the default confirming action.
        - **Drop**: Packets are dropped.
        - **Mark CoS**: Packets are marked by DiffServ with the specified CoS value before being forwarded. This selection requires that the Mark CoS field is set. You must select a CoS value from 0 to 7 from the menu.
        - **Mark IP Precedence**: These packets are marked by DiffServ with the specified IP Precedence value before being forwarded. This selection requires that the

Mark IP Precedence field is set. You must select an IP precedence value from 0 to 7 from the menu.

- **Mark IP DSCP**: Packets are marked by DiffServ with the specified DSCP value before being forwarded. This selection requires that the DSCP field is set. You must either select a DSCP code from the menu or enter an IP DSCP value from 0 to 63 in the field next to the menu. A value that you enter in the field overrides any selection from the menu. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header.

c. **Violate Action**: Specify a policy action for packets that violate the simple policy.

- **Send**: Packets are forwarded unmodified. This is the default violating action.
- **Drop**: Packets are dropped.
- **Mark CoS**: Packets are marked by DiffServ with the specified CoS value before being forwarded. This selection requires that the Mark CoS field is set. You must select a CoS value from 0 to 7 from the menu.
- **Mark IP Precedence**: These packets are marked by DiffServ with the specified IP Precedence value before being forwarded. This selection requires that the Mark IP Precedence field is set. You must select an IP precedence value from 0 to 7 from the menu.
- **Mark IP DSCP**: Packets are marked by DiffServ with the specified DSCP value before being forwarded. This selection requires that the DSCP field is set. You must either select a DSCP code from the menu or enter an IP DSCP value from 0 to 63 in the field next to the menu. A value that you enter in the field overrides any selection from the menu. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header.

○ **Two Rate**: Select this radio button to define a two-rate traffic policing style for the class. A two-rate policy differs from a simple policy in that the two-rate policy allows you to define a peak rate and a peak burst size, in addition to the committed rate and committed burst size. A two-rate policy supports two data rates and results in and action for packets that conform to the policy and an action for packets that violate the policy. By default, the two-rate policy is color blind, and color classes do not apply.

a. Specify the traffic rates and bust sizes for packets that conform to the two-rate policy.

- **Committed Rate**: From the menu, select **Percent** or **Kbps**, and then specify the committed rate that is applied to conforming packets by entering a value in the range from 1 to 100 percent or from 1 to 4294967295 kilobits-per-second (Kbps). The committed rate is used to guarantee the allowed bandwidth of conforming traffic.
- **Committed Burst Size**: Specify the burst size that is applied to conforming packets by entering a value in the range from 1 to 128 KBytes (KB). The committed burst size is used to limit the maximum number of bytes in a burst (a

very short instance) of conforming traffic.

- **Peak Rate**: Specify the peak rate that is applied to conforming packets by entering a value in the range from 1 to 4294967295 kilobits-per-second (Kbps). The peak rate is used to limit how much a burst of conforming traffic can exceed the committed rate.
- **Peak Burst Size**: Specify the peak burst size that is applied to conforming packets by entering a value in the range from 1 to 128 KBytes (KB). The peak burst size is used to limit how much a burst of conforming traffic can exceed the peak rate.

b. **Conform Action**: Specify a policy action for packets that conform to the two-rate policy.
   - **Send**: Packets are forwarded unmodified. This is the default confirming action.
   - **Drop**: Packets are dropped.
   - **Mark CoS**: Packets are marked by DiffServ with the specified CoS value before being forwarded. This selection requires that the Mark CoS field is set. You must select a CoS value from 0 to 7 from the menu.
   - **Mark IP Precedence**: These packets are marked by DiffServ with the specified IP Precedence value before being forwarded. This selection requires that the Mark IP Precedence field is set. You must select an IP precedence value from 0 to 7 from the menu.
   - **Mark IP DSCP**: Packets are marked by DiffServ with the specified DSCP value before being forwarded. This selection requires that the DSCP field is set. You must either select a DSCP code from the menu or enter an IP DSCP value from 0 to 63 in the field next to the menu. A value that you enter in the field overrides any selection from the menu. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header.

c. **Exceed Action**: Specify a policy action for packets that exceed the two-rate policy.
   - **Send**: Packets are forwarded unmodified. This is the default violating action.
   - **Drop**: Packets are dropped.
   - **Mark CoS**: Packets are marked by DiffServ with the specified CoS value before being forwarded. This selection requires that the Mark CoS field is set. You must select a CoS value from 0 to 7 from the menu.
   - **Mark IP Precedence**: These packets are marked by DiffServ with the specified IP Precedence value before being forwarded. This selection requires that the Mark IP Precedence field is set. You must select an IP precedence value from 0 to 7 from the menu.
   - **Mark IP DSCP**: Packets are marked by DiffServ with the specified DSCP value before being forwarded. This selection requires that the DSCP field is set. You must either select a DSCP code from the menu or enter an IP DSCP value from 0 to 63 in the field next to the menu. A value that you enter in the field overrides

any selection from the menu. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header.

d. **Violate Action**: Specify a policy action for packets that violate the two-rate policy.
   - **Send**: Packets are forwarded unmodified. This is the default violating action.
   - **Drop**: Packets are dropped.
   - **Mark CoS**: Packets are marked by DiffServ with the specified CoS value before being forwarded. This selection requires that the Mark CoS field is set. You must select a CoS value from 0 to 7 from the menu.
   - **Mark IP Precedence**: These packets are marked by DiffServ with the specified IP Precedence value before being forwarded. This selection requires that the Mark IP Precedence field is set. You must select an IP precedence value from 0 to 7 from the menu.
   - **Mark IP DSCP**: Packets are marked by DiffServ with the specified DSCP value before being forwarded. This selection requires that the DSCP field is set. You must either select a DSCP code from the menu or enter an IP DSCP value from 0 to 63 in the field next to the menu. A value that you enter in the field overrides any selection from the menu. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header.

9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.
    The following table describes the view-only fields on the page.

Table 167.DiffServ policy configuration, policy attributes

Field	Description
Policy Name	The name of the DiffServ policy
Policy Type	The type of the policy (In or Out)
Member Class Name	The class instances that are associated within the policy

Rename an existing DiffServ policy

You can change the name of an existing DiffServ policy.

To rename an existing DiffServ policy:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **QoS > DiffServ > Advanced > Policy Configuration**.
   The Policy Configuration page displays.
3. Select the check box next to the policy name.
4. In the **Policy Name** field, specify the new name.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Change the policy attributes for an existing DiffServ policy

You can change the policy attributes for an existing DiffServ policy.

To change the policy attributes for an existing DiffServ policy:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > Policy Configuration**.
   The Policy Configuration page displays.
3. Click the policy name, which is a hyperlink.
   The page on which you can change the policy attributes displays.
4. Change the policy attributes as needed.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Change or remove a class from an existing DiffServ policy

You can change a class in an existing DiffServ policy or remove a class from the policy.

To change or remove a class from an existing DiffServ policy:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > Policy Configuration**.
   The Policy Configuration page displays.
3. Select the check box next to the policy name.

4. Do one of the following:
   - To change the class, select another class from the **Member Class** menu.
   - To remove the class, select **None**, from the **Member Class** menu.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete a DiffServ policy

You can delete a DiffServ policy that you no longer need.

To delete a DiffServ policy:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > Policy Configuration**.
   The Policy Configuration page displays.
3. Select the check box next to the policy name.
4. Click the **Delete** button.
   The policy is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the DiffServ service interface

You can assign (attach) both an inbound policy and an outbound policy to an interface.

Attach DiffServ policies to an interface

You can attach an inbound policy, outbound policy, or both types of policies to an interface.

To attach DiffServ policies to an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > Service Interface Configuration**.
   The Service Interface Configuration page displays.

3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Policy In Name** menu, select the name of a policy for inbound traffic.
6. From the **Policy Out Name** menu, select the name of a policy for outbound traffic.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 168.Service Interface Configuration information

Field	Description
Direction	The traffic direction (In, Out, or both In and Out) of the policy or policies on the service interface.
Operational Status	The operational status of this service interface (either Up or Down). The operational status is shown as Up if all of the following conditions are true: • The attached class is valid and includes at least one matching rule. • The attached policy is valid and includes at least one attribute. • The port is enabled, that is, the physical link of the port is in the *up* state.

Change one or both DiffServ policies for an interface

You can change the inbound policy, outbound policy, or both policies for an interface.

To change one or both DiffServ policies for an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > Service Interface Configuration**.
   The Service Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button
5. To select another inbound policy, from the **Policy In Name** menu, select another policy name.
6. To select another outbound policy, from the **Policy Out Name** menu, select another policy name.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Remove one or both DiffServ policies from an interface

You can remove the inbound policy, outbound policy, or both policies from an interface.

To remove one or both DiffServ policies from an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **QoS > DiffServ > Advanced > Service Interface Configuration**.
   The Service Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button
5. To remove the inbound policy, from the **Policy In Name** menu, select None.
6. To remove the outbound policy, from the **Policy Out Name** menu, select None.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Display DiffServ service statistics

You can display service-level statistical information about all interfaces to which DiffServ policies are attached.

To display DiffServ service statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **QoS > DiffServ > Advanced > Service Statistics**.
   The Service Statistics page displays.

3. Use the **Search** menu to search for DiffServ statistics by interface or member class:
   - To search by interface, select **Interface**, enter the interface in unit/slot/port format (for example, 1/0/13), and click the **Go** button. If the entry exists, the entry is displayed as the first entry, followed by the remaining entries.
   - To search by member class, select **Member Class**, enter the member class, and click the **Go** button. If an entry with a matching member class exists, the entry is displayed as the first entry, followed by the remaining entries. An exact match is required.
4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 169.DiffServ Service Statistics information

Field	Description
Interface	The interface with a DiffServ policy attached.
Direction	The direction (In or Out) for which the DiffServ policy is attached.
Policy Name	The name of the DiffServ policy. If more than one class is attached to the policy, each policy instance (with the same policy name) is displayed.
Operational Status	The operational status of the DiffServ policy (Up or Down). The status is displayed as Up if all of the following conditions are true: ○ The attached class is valid and includes at least one matching rule. ○ The attached policy is valid and includes at least one attribute. ○ The port is enabled, that is, the physical link of the port is in the *up* state.
Member Classes	All DiffServ classes that are defined as members of the selected policy name. Select a member class name to display its statistics. If no class is associated with the selected policy, then the list is empty.
Offered Packets	The total number of packets offered to all class instances on this service policy before their defined DiffServ policy treatment is applied. This is the total number for the interface and direction.
Discarded Packets	The total number of packets discarded for any reason for all class instances on this service policy when the DiffServ policy treatment is applied. This is the total number for the interface and direction.

Manage Switch Security

The chapter covers the following topics:

- User accounts and passwords
- RADIUS servers
- TACACS+ servers
- Authentication lists
- Current login sessions
- HTTP and HTTPS management access
- SSH management access
- Telnet management access
- Console port management access
- Denial of service
- Management access profiles and rules
- Port authentication
- MAC filters for traffic control
- Port security
- Private port groups
- Protect ports
- Private VLANs
- Storm control
- DHCP snooping
- DHCPv6 snooping
- IP source guard interfaces
- IPv6 source guard interfaces
- Dynamic ARP inspection
- Captive portals
- Access control lists

User accounts and passwords

You can configure user accounts and login passwords. By default, two user accounts exist:

- admin, with read and write (read/write) privileges
- guest, with read-only privileges

The account names are not case-sensitive.

The first time that you log in as an admin user to the main UI, no password is required (that is, the password is blank). After you log in for the first time, you are required to specify a local device password that you then must use to log in again. Each subsequent time that you log in, you must use your local device password.

A guest user cannot log in until the admin user specifies a password for the guest user.

If you log in as an admin user with read/write privileges, you can assign passwords, set security

parameters for the default accounts, and add and delete accounts (other than the admin account), up to a maximum of six accounts. As an admin user with read/write privileges, you modify all data on the main UI pages.

Add or change a user account

You can add or change a user account, including the predefined users accounts *admin* and *guest*.

To add or change a user account:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Local User > User Management**.
   The User Management page displays.
3. If you change the settings for an existing user account, select the check box for the user account.
   You cannot change the name of an existing user account.
4. If you are adding a new user account, In the **User Name** field, enter a name.
   You can enter a user name only when you are creating an account. User names are up to 64 characters in length and are not case-sensitive. Valid characters include all the alphanumeric characters as well as the hyphen (-) and underscore (_) characters. The user name default is not valid. User names once created cannot be changed or modified.
5. From the **Edit Password** menu, select **Enable**.
   By default, the selection is Disable. Enable this option only if you are defining or changing the password options.
6. In the **Password** field, enter the password for the account.
   The characters do not display as they are typed; only asterisks (*) show. Passwords are up to eight alphanumeric characters in length, and are case-sensitive.
7. In the **Confirm Password** field, enter the password again, to confirm that you entered it correctly.
   This field does not display the password as it is typed, but shows asterisks (*).
8. From the **Access Mode** menu, select **Read_Write** or **Read_Only**.
9. From the **Encryption Type** menu, select the encryption strength:
   - **SHA256**: A secure hash algorithm that is compatible with many devices.
   - **SHA512**: A secure hash algorithm that is superior to SHA256. This is the default type.

10. Form the **Multifactor Authentication Mode** menu, select to enable or disable multifactor authentication:
    - **Disable**: Multifactor authentication is disabled. This is the default selection.
    - **Enable**: Multifactor authentication is enabled and an authentication email is sent to the email address that you must enter in the **Multifactor Authentication Email Information** field.
11. Do one of the following:
    - If you are adding a new user account, click the **Add** button. Your settings are saved. The user account is added.
    - If you are changing the settings for a user account, click the **Apply** button. Your settings are saved.
12. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 170.User management information

Field	Description
Lockout Status	Indicates if the user account is locked (True or False)
Password Expiration Date	The password expiration date, if any

Delete a user account

You can delete a user account that is no longer required.

To delete a user account:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Local User > User Management**.
   The User Management page displays.
3. Select the check box for the user.
4. Click the **Delete** button.
   The user is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure user password requirements

To configure user password requirements:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Local User > User Password Configuration**.
   The User Password Configuration page displays.
3. In the **Password Minimum Length** field, type the minimum character length for any new local user password.
   The default is eight characters. The range is from 0 to 64 characters.
4. In the **Password Aging (days)** field, type the maximum number of days during which a user password is valid, counted from the day that the password is defined.
   After a password expires, the user must define a new password when logging in. The default is 0, which means that passwords do not expire. The range is from 0 to 365 days.
5. In the **Password History** field, type the number of previous passwords that the switch stores to prevent users from reusing passwords.
   The default is 0, which means that the switch does not store previous passwords. The range is from 0 to 10 passwords.
6. In the **Lockout Attempts** field, specify the number of times a user is allowed to fail a local authentication attempt before the user account is locked.
   The default is 5 times. If you enter 0, a failed local authentication attempt does not cause the user account to be locked. The range is from 0 to 5 times.
7. In the **Unlock Time** field, specify the time in minutes after which a locked user account is unlocked.
   The default is 5 minutes. The range is from 1 to 60 minutes.

8. Select an Unlock Timer Mode radio button to enable or disable the locking of a user account after too many failed login attempts:
   - **Disable**: A user account is never locked.
   - **Enable**: A user account is locked for a period of time after too many failed login attempts. After you enter the wrong password three times, the switch blocks further attempts to log in. The block lasts longer each time until you get the password right:
     - After 3 times: 5 minutes
     - After 6 times: 10 minutes
     - After 9 times: 20 minutes
     - After 12 times: 40 minutes
     - After 15 times: 60 minutes
     - After 18 times: 60 minutes
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

Enable multi-factor authentication on the switch

If you enable multi-factor authentication (MFA) on the switch, it applies to the following types of access:

- Main UI
- CLI access over Telnet
- CLI access over SSH

MFA does not apply to console access.
You can set up MFA for an individual user (see Add or change a user account on page 556) and configure the email address that is used for one time passwords (OTPs), but in order for MFA to take effect, MFA must also be globally enabled on the switch.
For information about how to use the CLI to configure a mail server on the switch, see the CLI manual, which you download by visiting netgear.com/support/enterprise/.

To enable multifactor authentication on the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Local User > MFA Global Configuration**.
   The MFA Global Configuration page displays.

3. Select a Multifactor Authentication Mode radio button to enable or disable multifactor authentication:
   - **Disable**: Multifactor authentication is disabled on the switch. This is the default setting.
   - **Enable**: Multifactor authentication is enabled on the switch.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Change the privileged EXEC CLI password

You can change the privileged EXEC password for access to the privileged EXEC mode in the command-line interface (CLI). Passwords are a maximum of 64 alphanumeric characters. The password is case-sensitive. You can also change the encryption strength.

To change the privileged EXEC CLI password for use in the CLI and the encryption strength:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Enable Password**.
   The Enable Password page displays.
3. From the **Encryption Type** menu, select the encryption strength:
   - **SHA256**: A secure hash algorithm that is compatible with many devices.
   - **SHA512**: A secure hash algorithm that is superior to SHA256. This is the default type.
4. In the **Password** field, type the password.
   The password can be a maximum of 64 alphanumeric characters.
5. In the **Confirm Password** field, type the password again, to confirm that you entered it correctly.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Change the console, Telnet, or SSH password

A line password is a password that is used for console, Telnet, or SSH access.
You can change a password and the encryption strength.

To configure the console, Telnet, or SSH password and the encryption strength:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Line Password**.
   The Line Password page displays.
3. To change the console password and encryption strength, do the following:
   a. Select the **Console** check box.
   b. From the **Encryption Type** menu, select the encryption strength:
      - **SHA256**: A secure hash algorithm that is compatible with many devices.
      - **SHA512**: A secure hash algorithm that is superior to SHA256. This is the default type.
   c. In the **Console Password** field, type the console password. The password can be a maximum of 64 alphanumeric characters.
   d. In the **Confirm Console Password** field, type the password again to confirm that you typed it correctly.
4. To change the Telnet password and encryption strength, do the following:
   a. Select the **Telnet** check box.
   b. From the **Encryption Type** menu, select the encryption strength:
      - **SHA256**: A secure hash algorithm that is compatible with many devices.
      - **SHA512**: A secure hash algorithm that is superior to SHA256. This is the default type.
   c. In the **Console Password** field, type the console password. The password can be a maximum of 64 alphanumeric characters.
   d. In the **Confirm Console Password** field, type the password again to confirm that you typed it correctly.
5. To change the SSH password and encryption strength, do the following:
   a. Select the **SSH** check box.
   b. From the **Encryption Type** menu, select the encryption strength:
      - **SHA256**: A secure hash algorithm that is compatible with many devices.
      - **SHA512**: A secure hash algorithm that is superior to SHA256. This is the default type.
   c. In the **Console Password** field, type the console password. The password can be a maximum of 64 alphanumeric characters.
   d. In the **Confirm Console Password** field, type the password again to confirm that you typed it correctly.
6. Click the **Apply** button.
   Your settings are saved.

7. To save the settings to the running configuration, click the **Save** icon.

RADIUS servers

RADIUS servers provide additional security for networks. The RADIUS server maintains a user database, which contains per-user authentication information. The switch passes information to the configured RADIUS server, which can authenticate a user name and password before authorizing use of the network. RADIUS servers provide a centralized authentication method for the following:

- Web access
- Port access control (802.1X)

Configure the global RADIUS server settings

You can add information about multiple RADIUS servers on the network.
If you configure multiple RADIUS servers, consider the maximum delay time when you specify the maximum number of retransmissions (that is, the value that you enter in the **Max Number of Retransmits** field in the following procedure) and the time-out period (that is, the value that you enter in the **Timeout Duration** field in the following procedure) for RADIUS:

- For one RADIUS server, a retransmission does not occur until the configured time-out period expires without a response from the RADIUS server. In addition. the maximum number of retransmissions for one RADIUS server must pass before the switch attempts the next RADIUS server.
- Therefore, the maximum delay in receiving a RADIUS response on the switch equals the maximum number of retransmissions multiplied by the time-out period multiplied by the number of configured RADIUS servers. If the RADIUS request was generated by a user login attempt, all user interfaces are blocked until the switch receives a RADIUS response.

To configure the global RADIUS server settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > RADIUS > Global Configuration**.
   The Global Configuration page displays.
   The Current Server Address field is blank if no servers are configured (see Configure a RADIUS authentication server on the switch on page 564).
   The switch supports up to three configured RADIUS servers. If more than one RADIUS server is configured, the current server is the primary server. If no servers is configured as the primary server, the current server is the most recently added RADIUS server.

3. From the **Source Interface** menu, select the source interface that the switch must use to communicate with the RADIUS servers. By default, the following options display in the menu:
   - **None**: The primary IP address of the originating (outbound) interface is used as the source address.
   - **VLAN 1**: The primary IP address of VLAN 1 is used as the source address. This is the default selection.
   - **Service Port**: The management port IP address is used as the source address.

   Depending on the configuration of your switch, the following options can display:
   - **Another VLAN ID**: The primary IP address of a VLAN other than VLAN 1 is used as the source address.
   - **Routing interface**: The primary IP address of a routing interface is used as the source address.
   - **Routing VLAN**: The primary IP address of a VLAN routing interface is used as the source address.
   - **Routing loopback interface**: The primary IP address of a routing loopback interface is used as the source address.
   - **Different**: For some features, *Different* can display. This means that the source interface is configured separately.
4. In the **Max Number of Retransmits** field, specify the maximum number of times a request packet is retransmitted to the RADIUS server.
   The range is from 1 to 15. The default value is 4.

   CAUTION: The maximum delay in receiving a RADIUS response on the switch equals the maximum number of retransmissions multiplied by the time-out period multiplied by the number of configured RADIUS servers. If the RADIUS request was generated by a user login attempt, all user interfaces are blocked until the switch receives a RADIUS response.
5. In the **Timeout Duration** field, specify the time-out value, in seconds, for request retransmissions.
   The range is from 1 to 30. The default value is 5.
6. Select the Accounting Mode **Disable** or **Enable** radio button.
   This selection specifies if the RADIUS accounting mode is enabled on the current server. By default, the selection is Disable.
7. Select the RADIUS Attribute 4 Mode **Disable** or **Enable** radio button.
   This selection specifies if RADIUS attribute 4 is enabled on the current server. By default, the selection is Disable. If you select the **Enable** radio button, the **RADIUS Attribute 4 Value** field displays, enabling you to enter an IPv4 address for the RADIUS attribute 4 option.

8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 171.Radius configuration information

Field	Description
Current Server Address	The address of the current RADIUS server. This field is blank if no servers are configured.
Number of Configured Authentication Servers	The number of configured authentication RADIUS servers. The value can range from 0 to 32.
Number of Configured Accounting Servers	The number of RADIUS accounting servers configured. The value can range from 0 to 32.
Number of Named Authentication Server Groups	The number of named RADIUS server authentication groups configured.
Number of Named Accounting Server Groups	The number of named RADIUS server accounting groups configured.

Configure a RADIUS authentication server on the switch

You can configure and display various settings for communication between the switch and a RADIUS authentication server.

Add a RADIUS authentication server to the switch

You can add information about a RADIUS authentication server to the switch and display or reset the RADIUS authentication server statistics.

To add a RADIUS authentication server to the switch and display or reset the RADIUS authentication server statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > RADIUS > Server Configuration**.
   The Server Configuration page displays.
3. In the **Radius Server IP Address** field, specify the IP address of the RADIUS server.

4. In the **Radius Server Name** field, specify a name for the RADIUS server.
   This name helps you to identify the RADIUS server.
5. In the **Port** field, specify the UDP port number of the RADIUS authentication server.
   The range is from 1 to 65535.
6. From the **Secret Configured** menu, select **Yes**.
   You must select **Yes** before you can configure the RADIUS secret. After you add the RADIUS server, this field indicates whether the shared secret for this server was configured.
7. In the **Secret** field, type the shared secret text string used for authenticating and encrypting all RADIUS communications between the switch and the RADIUS server.
   This secret must match the one that is configured on the RADIUS server.
8. From the **Primary Server** menu, select one of the following options:
   - **Yes**: The server functions as the primary RADIUS authentication server.
   - **No**: The server functions as a secondary RADIUS authentication server.
9. From the **Message Authenticator** menu, select **Enable** or **Disable** to specify whether the message authenticator attribute for the selected server is enabled.
   The message authenticator adds protection to RADIUS messages by using an MD5 hash to encrypt each message. The shared secret is used as the key, and if the message fails to be verified by the RADIUS server, it is discarded.
10. From the **Server Type** menu, select one of the following options:
    - **Netgear**: A NETGEAR-specific RADIUS server.
    - **Standard**: A third-party RADIUS server.
11. Click the **Add** button.
    The server is added to the switch.
12. To reset the authentication server and RADIUS statistics to their default values, click the **Clear Counters** button.
13. To save the settings to the running configuration, click the **Save** icon.

The Current field shows if the server is currently in use as the RADIUS authentication server. The following table describes the view-only fields on the page.

Table 172.RADIUS authentication server statistics information

Field	Description
Radius Server	The address of the RADIUS server or the name of the RADIUS server for which the statistics are displayed
Round Trip Time	The time interval, in hundredths of a second, between the most recent access-reply/access-challenge and the access-request that matched it from this RADIUS authentication server
Access Requests	The number of RADIUS access-request packets sent to this server. This number does not include retransmissions.
Access Retransmissions	The number of RADIUS access-request packets retransmitted to this server
Access Accepts	The number of RADIUS access-accept packets, including both valid and invalid packets, that were received from this server
Access Rejects	The number of RADIUS access-reject packets, including both valid and invalid packets, that were received from this server
Access Challenges	The number of RADIUS access-challenge packets, including both valid and invalid packets, that were received from this server
Malformed Access Responses	The number of malformed RADIUS access-response packets received from this server. Malformed packets include packets with an invalid length. Bad authenticators or signature attributes or unknown types are not included in malformed access-responses.
Bad Authenticators	The number of RADIUS access-response packets containing invalid authenticators or signature attributes received from this server
Pending Requests	The number of RADIUS access-request packets destined for this server that did not yet time out or receive a response
Timeouts	The number of authentication time-outs to this server
Unknown Types	The number of RADIUS packets of unknown type that were received from this server on the authentication port
Packets Dropped	The number of RADIUS packets received from this server on the authentication port and dropped for some other reason

Modify the settings for a RADIUS authentication server on the switch

You can modify the settings for a RADIUS authentication server on the switch.

To modify the settings for a RADIUS authentication server on the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > RADIUS > Server Configuration**.
   The Server Configuration page displays.
3. Select the check box next to the server IP address.
4. Modify the configuration for the selected server.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Remove a RADIUS authentication server from the switch

If you no longer need a RADIUS authentication server, you can remove it from the switch.

To remove a RADIUS authentication server from the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > RADIUS > Server Configuration**.
   The Server Configuration page displays.
3. Select the check box next to the IP address of the server.
4. Click the **Delete** button.
   The RADIUS server is removed.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Configure a RADIUS accounting server on the switch

You can configure and display various settings for communication between the switch and a RADIUS accounting server.

Add a RADIUS accounting server to the switch

You can add information about a RADIUS accounting server to the switch and display or reset the RADIUS accounting server statistics.

To add a RADIUS accounting server to the switch and display or clear the RADIUS accounting server statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > RADIUS > Accounting Server Configuration**.
   The Accounting Server Configuration page displays.
3. In the **Accounting Server IP Address** field, specify the IP address of the RADIUS accounting server.
4. In the **Accounting Server Name** field, enter the name of the accounting server.
5. In the **Port** field, specify the UDP port number of the RADIUS accounting server.
   The range is from 1 to 65535.
6. From the **Secret Configured** menu, select **Yes** to add a RADIUS secret in the next field.
   After you add the RADIUS accounting server, this field indicates whether the shared secret for this server is configured.
7. In the **Secret** field, type the shared secret to use with the specified accounting server.
8. From the **Accounting Mode** menu, select to enable or disable the RADIUS accounting mode:
   - **Enable**: All configured RADIUS accounting servers are enabled on the switch.
   - **Disable**: All configured RADIUS accounting servers are disabled on the switch.

   NOTE: Your selection from the **Accounting Mode** menu applies to all configured accounting servers, not only to the server that you are adding.
9. Click the **Add** button.
   The server is added to the switch.
10. To reset the authentication server and RADIUS statistics to their default values, click the **Clear Counters** button.
11. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 173.RADIUS accounting server statistics information

Field	Description
Accounting Server	The accounting server associated with the statistics
Round Trip Time	The time interval, in hundredths of a second, between the most recent accounting-response and the accounting-request that matched it from this RADIUS accounting server
Accounting Requests	The number of RADIUS accounting-request packets sent not including retransmissions
Accounting Retransmissions	The number of RADIUS accounting-request packets retransmitted to this RADIUS accounting server
Accounting Responses	The number of RADIUS packets received on the accounting port from this server
Malformed Accounting Responses	The number of malformed RADIUS accounting-response packets received from this server. Malformed packets include packets with an invalid length. Bad authenticators and unknown types are not included as malformed accounting responses.
Bad Authenticators	The number of RADIUS accounting-response packets that contained invalid authenticators received from this accounting server
Pending Requests	The number of RADIUS accounting-request packets sent to this server that did not yet time out or receive a response
Timeouts	The number of accounting time-outs to this server
Unknown Types	The number of RADIUS packets of unknown type that were received from this server on the accounting port
Packets Dropped	The number of RADIUS packets that were received from this server on the accounting port and dropped for some other reason

Modify the settings for a RADIUS accounting server on the switch

You can modify the settings for a RADIUS accounting server on the switch.

To modify the settings for a RADIUS accounting server on the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Management Security > RADIUS > Accounting Server Configuration**.
   The Accounting Server Configuration page displays.
3. Select the check box next to the server IP address.
4. Modify the configuration for the accounting server.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Remove a RADIUS accounting server from the switch

If you no longer need a RADIUS accounting server, you can remove it from the switch.

To remove a RADIUS accounting server from the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > RADIUS > Accounting Server Configuration**.
   The Accounting Server Configuration page displays.
3. Select the check box next to the server IP address.
4. Click the **Delete** button.
   Your settings are saved. The RADIUS accounting server is removed.
5. To save the settings to the running configuration, click the **Save** icon.

TACACS+ servers

TACACS+ provides a centralized user management system, while still retaining consistency with RADIUS and other authentication processes. TACACS+ provides the following services:

- **Authentication**:Provides authentication during login and through user names and user-defined passwords.
- **Authorization**:Performed at login. When the authentication session is completed, an authorization session starts using the authenticated user name. The TACACS+ server checks the user privileges.

The TACACS+ protocol ensures network security through encrypted protocol exchanges between the device and TACACS+ server.

Configure the global TACACS+ settings

You can configure the global TACACS+ settings for communication between the switch and a TACACS+ server.

To configure the global TACACS+ settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > TACACS > TACACS+ Configuration**.
   The TACACS+ Configuration page displays.
3. In the **Key String** field, specify the authentication and encryption key for TACACS+ communications between the switch and the TACACS+ server.
   The range is from 0 to 128. The key must match the key configured on the TACACS+ server.
4. In the **Connection Timeout** field, specify the maximum number of seconds allowed to establish a TCP connection between the switch and the TACACS+ server.
   The range is from 1 to 30 seconds. The default is 5 seconds.
5. From the **Source Interface** menu, select the source interface that must be used for TACACS+ communication. By default, the following options display in the menu:
   - **None**: The primary IP address of the originating (outbound) interface is used as the source address.
   - **VLAN 1**: The primary IP address of VLAN 1 is used as the source address. This is the default selection.
   - **Service Port**: The management port IP address is used as the source address.

   Depending on the configuration of your switch, the following options can display:
   - **Another VLAN ID**: The primary IP address of a VLAN other than VLAN 1 is used as the source address.
   - **Routing interface**: The primary IP address of a routing interface is used as the source address.
   - **Routing VLAN**: The primary IP address of a VLAN routing interface is used as the source address.
   - **Routing loopback interface**: The primary IP address of a routing loopback interface is used as the source address.
   - **Different**: For some features, *Different* can display. This means that the source interface is configured separately.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Add a TACACS+ server to the switch

You can add up to five TACACS+ servers with which the switch can communicate.

To add a TACACS+ server to the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > TACACS > TACACS+ Server Configuration**.
   The TACACS+ Server Configuration page displays.
3. In the **TACACS+ Server** field, enter the TACACS+ server IP address.
4. In the **Priority** field, specify the priority for the TACACS+ server.
   The priority determines the order in which the TACACS+ servers are contacted when attempting to authenticate a user. A value of 0 is the highest priority. The range is from 0 to 65535.
5. In the **Port** field, specify the authentication port number for TACAS+ server sessions.
   The value must be in the range from 0 to 65535.
6. In the **Key String** field, specify the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server.
   The key can be from 0 to 128 characters. The key must match the key used on the TACACS+ server.
7. In the **Connection Timeout** field, specify the time that passes before the connection between the device and the TACACS+ server times out.
   The range is from 1 to 30 seconds.
8. Click the **Add** button.
   The server is added to the switch.
9. To save the settings to the running configuration, click the **Save** icon.

Modify the settings for a TACACS+ server on the switch

To modify the settings for a TACACS+ server on the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > TACACS+ > TACACS+ Server Configuration**.
   The TACACS+ Server Configuration page displays.

3. Select the check box next to the server IP address.
4. Modify the configuration for the selected TACACS+ server.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Remove a TACACS+ server from the switch

To remove a TACACS+ server from the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > TACACS+ > TACACS+ Server Configuration**.
   The TACACS+ Server Configuration page displays.
3. Select the check box next to the server IP address.
4. Click the **Delete** button.
   The TACACS+ server is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Authentication lists

An authentication login list specifies one or more authentication methods to validate switch or port access for an admin or guest user. Access can be HTTP or HTTPS access to the main UI or AV UI, or line access (that is, Telnet, SSH, or console access) to the CLI.

Configure a login authentication list

A login list specifies the authentication methods to be used to validate switch or port access. An authentication list can apply to users that access the switch through one of the following methods:

- **Secure Shell (SSH)**: For information about selecting an authentication list for SSH access, see SSH management access on page 588.
- **Telnet**: For information about selecting an authentication list for Telnet access, see Telnet management access on page 593.
- **Console port**: For information about selecting an authentication list for console port access, see Console port management access on page 596.

Two default lists are present: defaultList and networkList. You cannot delete these lists. By default, the preconfigured users (admin and guest) are assigned to both of these lists in the following way:

- For SSH and Telnet access, by default, the preconfigured users are assigned to the defaultList.
- For console access, by default, the preconfigured users are assigned to the networkList.

All newly added users are also assigned to the defaultList until you specifically assign them to a different list (see Configure the global 802.1X authentication settings on page 604).

To configure a login authentication list:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Authentication List > Login Authentication List**.
   The Login Authentication List page displays.
3. Do one of the following:
   - **Create a new list**: In the **List Name** field, specify a name. The name can be up to 15 alphanumeric characters long and is not case-sensitive.
   - **Modify an existing list**: Select the check box that is associated with the list.

4. From the menu in the **1** column, select the authentication method that must be used first in the selected authentication login list.
   If you select a method that does not time out as the first method, such as Local, no other method is tried, even if you specified more than one method. User authentication occurs in the order that you select the methods:
   - **Local**: The user's locally stored name and password are used for authentication. Because the Local method does not time out, if you select this option as the first method, no other method is tried, even if you specified more than one method. This is the default selection for method 1.
   - **Enable**: The privileged EXEC password is used for authentication.
   - **Line**: The line password is used for authentication.
   - **None**: The user is allowed access without authentication.
   - **Radius**: The user's name and password are authenticated using the RADIUS server instead of the local server.
   - **Tacacs**: The user's name and password are authenticated using the TACACS server instead of the local server.
5. From the menus in the **2**, **3**, **4**, **5**, and **6** columns, select the authentication methods, if any, that must be used in the selected authentication login list.
   If a previous method times out, the next method is used. For example, the authentication method that you select in the **2** column is tried after the authentication method that select in the **1** column in the previous step. Similarly, the authentication method that you select in the **3** column is tried after the authentication method that you select in the **2** column, and so on. If you select a method that does not time out, the next method is not tried.
   For the menus in the 2, 3, 4, 5, and 6 columns, in addition to the options described in the previous step, you can also select **N/A**. With the N/A option (which is the default selection) the authentication method is not used. This option is not available for the menu in the 1 column.
6. Do one of the following:
   - If you are adding a new list, click the **Add** button. Your settings are saved. The list is added to the table.
   - If you are changing an existing list, click the **Apply** button. Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Delete a login authentication list

You can delete a custom login authentication list that you no longer need. You cannot delete the defaultList and networkList.

To delete a login authentication list:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Authentication List > Login Authentication List**.
   The Login Authentication List page displays.
3. Select the check box that is associated with the list.
4. Click the **Delete** button.
   Your settings are saved. The list is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure an enable authentication list

An enable list specifies the authentication methods to validate privileged EXEC access in the CLI. An enable list can apply to users that access the switch CLI through one of the following methods:

- **Secure Shell (SSH)**: For information about selecting an enable list for SSH access, see SSH management access on page 588.
- **Telnet**: For information about selecting an enable list for Telnet access, see Telnet management access on page 593.
- **Console port**: For information about selecting an enable list for console port access, see Console port management access on page 596.

Two default lists are present: enableList and enableNetList. You cannot delete these lists.
The preconfigured users (admin and guest) are assigned to the enableList until you specifically select a different list.
All newly added users are also assigned to the enableList until you specifically select a different list.

To configure an enable authentication list:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Authentication List > Enable Authentication List**.
   The Enable Authentication List page displays.
3. Do one of the following:
   - **Create a new list**: In the List Name field, specify a name. The name can be up to 15 alphanumeric characters long and is not case-sensitive.
   - **Modify an existing list**: Select the check box that is associated with the list.

4. From the menu in the **1** column, select the authentication method that must be used first in the selected authentication login list.
   If you select a method that does not time out as the first method, such as Local, no other method is tried, even if you specified more than one method. User authentication occurs in the order that you select the methods:
   - **Enable**: The privileged EXEC password is used for authentication.
   - **Line**: The line password is used for authentication.
   - **None**: The user is allowed access without authentication.
   - **Radius**: The user's name and password are authenticated using the RADIUS server instead of the local server.
   - **Tacacs**: The user's name and password are authenticated using the TACACS server instead of the local server.
   - **Deny**: The user is always denied access.
5. From the menus in the **2**, **3**, **4**, and **5** columns, select the authentication methods, if any, that must be used in the selected authentication login list.
   If a previous method times out, the next method is used. For example, the authentication method that you select in the **2** column is tried after the authentication method that select in the **1** column in the previous step. Similarly, the authentication method that you select in the **3** column is tried after the authentication method that you select in the **2** column, and so on. If you select a method that does not time out, the next method is not tried.
   For the menus in the 2, 3, 4, and 5 columns, in addition to the options described in the previous step, you can also select **N/A**. With the N/A option (which is the default selection) the authentication method is not used. This option is not available for the menu in the 1 column.
6. Do one of the following:
   - If you are adding a new list, click the **Add** button. Your settings are saved. The list is added to the table.
   - If you are changing an existing list, click the **Apply** button. Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Delete an enable authentication list

You can delete a custom login authentication list that you no longer need. You cannot delete the enableList and enableNetList.

To delete an enable authentication list:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Management Security > Authentication List > Enable Authentication List**.
   The Enable Authentication List page displays.
3. Select the check box that is associated with the list.
4. Click the **Delete** button.
   Your settings are saved. The list is removed.

Configure the Dot1x authentication list

You can configure the dot1x authentication list, which specifies the authentication methods to validate port access for users associated with the dot1x list. Only a single dot1x list exists. For that list, you can select a single access method only. By default, access to ports does not require authentication, so no default access method exists.

To configure the dot1x authentication list:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Authentication List > Dot1x Authentication List**.
   The Dot1x Authentication List page displays.
3. Select the check box for the dot1x list.
4. Select the authentication method that must be used:
   - **Ias**: The user's ID and password in the internal authentication server (IAS) database is used for authentication.
   - **None**: The user is allowed access without authentication.
   - **Radius**: The user's name and password are authenticated using the RADIUS server instead of the local server.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Configure the HTTP authentication list

You can configure the HTTP authentication list, which specifies the authentication methods to validate switch or port access through HTTP, for example, through a web browser session. Only a single HTTP authentication list exists, which, by default, uses local authentication.

To configure the HTTP authentication list:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Authentication List > HTTP Authentication List**.
   The HTTP Authentication List page displays.
3. Select the check box for the HTTP list.
4. From the menu in the **1** column, select the authentication method that must be used first in the selected authentication login list.
   If you select a method that does not time out as the first method, such as Local, no other method is tried, even if you specified more than one method. User authentication occurs in the order that you select the methods:
   - **Local**: The user's locally stored name and password are used for authentication. Because the Local method does not time out, if you select this option as the first method, no other method is tried, even if you specified more than one method. This is the default selection for method 1.
   - **None**: The user is allowed access without authentication.
   - **Radius**: The user's name and password are authenticated using the RADIUS server instead of the local server.
   - **Tacacs**: The user's name and password are authenticated using the TACACS server instead of the local server.
5. From the menus in the **2**, **3**, and **4** columns, select the authentication methods, if any, that must be used in the selected authentication login list.
   If a previous method times out, the next method is used. For example, the authentication method that you select in the **2** column is tried after the authentication method that select in the **1** column in the previous step. Similarly, the authentication method that you select in the **3** column is tried after the authentication method that you select in the **2** column, and so on. If you select a method that does not time out, the next method is not tried.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Configure the HTTPS authentication List

You can configure the HTTPS authentication list, which specifies the authentication methods to validate switch or port access through secure HTTP (HTTPS), for example, through a web browser session. Only a single HTTPS authentication list exists, which, by default, uses local authentication.

To configure the HTTPS authentication list:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Authentication List > HTTPS Authentication List**.
   The HTTPS Authentication List page displays.
3. Select the check box for the HTTPS list.
4. From the menu in the **1** column, select the authentication method that must be used first in the selected authentication login list.
   If you select a method that does not time out as the first method, such as Local, no other method is tried, even if you specified more than one method. User authentication occurs in the order that you select the methods:
   - **Local**: The user's locally stored name and password are used for authentication. Because the Local method does not time out, if you select this option as the first method, no other method is tried, even if you specified more than one method. This is the default selection for method 1.
   - **None**: The user is allowed access without authentication.
   - **Radius**: The user's name and password are authenticated using the RADIUS server instead of the local server.
   - **Tacacs**: The user's name and password are authenticated using the TACACS server instead of the local server.
5. From the menus in the **2**, **3**, and **4** columns, select the authentication methods, if any, that must be used in the selected authentication login list.
   If a previous method times out, the next method is used. For example, the authentication method that you select in the **2** column is tried after the authentication method that select in the **1** column in the previous step. Similarly, the authentication method that you select in the **3** column is tried after the authentication method that you select in the **2** column, and so on. If you select a method that does not time out, the next method is not tried.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Current login sessions

You can display information about current Telnet, serial, SSH, HTTP, or HTTPS login sessions to the switch.

To display the current login sessions:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Management Security > Login Sessions**.
   The Login Sessions page displays.
   The following table describes the fields that are shown in the table.

Table 174.Login Sessions

Field	Description
ID	Identifies the ID of this row
User Name	The name of the user for whom the session is open
Connection From	The IP address of the computer from which the user is connected
Idle Time	The idle session time, if any
Session Time	The total session time
Session Type	The type of session: Telnet, Serial, SSH, HTTP, or HTTPS

HTTP and HTTPS management access

You can configure HTTP and secure HTTP (HTTPS) access to the switch's management interface.

Configure the HTTP access settings

You can configure the HTTP access settings on the switch. If you access the switch main UI, HTTP is the default access method.

To configure the HTTP access settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > HTTP > HTTP Configuration**.
   The HTTP Configuration page displays.

3. Select the HTTP Access **Disable** or **Enable** radio button:
   - **Disable**: You cannot access the switch main UI from an HTTP session over a web browser.
   - **Enable**: You can access the switch main UI from an HTTP session over a web browser. This is the default setting.
4. In the **HTTP Port** field, enter the HTTP port number.
   The valid range is 80 and from 1025 to 65535. The default port number is 80.
5. In the **HTTP Session Soft Timeout (Minutes)** field, specify the number of minutes an HTTP session can be idle before a time-out occurs.
   The value must be in the range from 0 to 60 minutes. The default value is 15 minutes.
   After the session is inactive and times out, you are automatically logged out and must reenter the password to access the main UI. A value of zero means that the session does not time out.
6. In the **HTTP Session Hard Timeout (Hours)** field, specify the hard time-out for HTTP sessions.
   This time-out is unaffected by the activity level of the session. The value must be in the range from 0 to 168 hours. A value of zero means that the session does not time out. The default value is 24 hours.
7. In the **Maximum Number of HTTP Sessions** field, specify the maximum number of HTTP sessions that can exist at the same time.
   The value must be in the range of 0 to 16. The default value is 16.
8. Click the **Apply** button.
   Your settings are saved.
   The Authentication List field displays HttpList, which is the default list for HTTP access.
9. To save the settings to the running configuration, click the **Save** icon.

Configure the HTTPS access settings

Secure HTTP (HTTPS) enables the transmission of HTTP over an encrypted Secure Sockets Layer (SSL) or Transport Layer Security (TLS) connection. When you manage the switch over the main UI, HTTPS can help ensure that communication between the management system and the switch is protected from eavesdroppers and man-in-the-middle attacks.

To configure the HTTPS access settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Access > HTTPS > HTTPS Configuration**.
   The HTTPS Configuration page displays.
3. Select the Admin Mode **Disable** or **Enable** radio button:
   - **Disable**: You cannot access the switch main UI from an HTTPS session over a web browser.

     NOTE: You can download SSL certificates only when HTTPS is disabled.

   - **Enable**: You can access the switch main UI from an HTTPS session over a web browser. This is the default setting.

   The Operational Status field shows if HTTPS is operationally enabled or disabled. For example, if all certificates expire, HTTPS is operationally disabled.
   The TLS Version field shows the TLS version, which is always 1.2.
4. In the **HTTPS Port** field, specify the HTTPS port number.
   The valid range is from 1025 to 65535. The default port number is 443.
5. In the **HTTPS Session Soft Timeout (Minutes)** field, specify the number of minutes an HTTPS session can be idle before a time-out occurs.
   The value must be in the range from 0 to 60 minutes. The default value is 15 minutes.
   After the session is inactive and times out, you are automatically logged out and must reenter the password to access the main UI. A value of zero means that the session does not time out.
6. In the **HTTPS Session Hard Timeout (Hours)** field, specify the hard time-out for HTTPS sessions.
   This time-out is unaffected by the activity level of the session. The value must be in the range from 0 to 168 hours. A value of zero means that the session does not time out. The default value is 24 hours.
7. In the **Maximum Number of HTTPS Sessions** field, specify the maximum number of HTTPS sessions that can exist at the same time.
   The value must be in the range of 0 to 16. The default value is 16.
8. Click the **Apply** button.
   Your settings are saved.
   The TLS Version field display 1.2, which is the default version for the switch.
   The Authentication List field displays HttpsList, which is the default list for HTTPS access.
9. To save the settings to the running configuration, click the **Save** icon.

Browser security message with HTTPS access

After you enable HTTPS access (see Configure the HTTP access settings on page 581) and you

attempt to access the main UI, your browser might display a security warning because of the self-signed certificate on the switch. This is expected behavior. You can proceed, or add an exception for the security warning.
To proceed with a security warning or add an exception for a security warning:

- **Google Chrome**: Click the **ADVANCED** link. Then, click the **Proceed to x.x.x.x (unsafe)** link, in which **x.x.x.x** represents the IP address of the switch.
- **Apple Safari**: Click the **Show Details** button. Then, click the **visit this website** link. If a warning pop-up window displays, click the **Visit Website** button. If another pop-up window display to let you confirm changes to your certificate trust settings, enter your Mac user name and password and click the **Update Setting** button.
- **Mozilla Firefox**: Click the **ADVANCED** button. Then, click the **Add Exception** button. In the pop-up window that displays, click the **Confirm Security Exception** button.
- **Microsoft Edge**: Select **Details > Go on to the webpage**.
- **Microsoft Internet Explorer**: Click the **Continue to this website (not recommended)** link.

NOTE: For information about installing a specific security certificate on the switch, see Manage certificates for HTTPS access on page 584 and Download and install an SSL security certificate file on the switch on page 779.

Manage certificates for HTTPS access

You can manage SSL certificates for HTTPS access. The switch supports two certificates, one of which one can be the active certificate.
You can either generate a certificate on the switch or transfer a certificate from a server to the switch.

Display the status of the SSL certificates

The switch can contain self-generated certificates that you added and certificates that you transferred to the switch.

To display the status of the SSL certificates:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > HTTPS > Certificate Management**.
   The Certificate Management page displays.

3. From the **Certificate Number** menu, select the certificate, that is, select **1** or **2**.

> **NOTE:** A certificate that you transfer (upload) to the switch is assigned number 1.

The Certificate Present field displays if the selected certificate is present on the switch.
The Certificate Expired field displays if the selected certificate is expired.
The Certificate Status section displays the following for all certificates on the switch:

- Active Certificate: The active certificates on the switch. If HTTPS is disabled, None is displayed. Otherwise, the certificate numbers are displayed.
- Expired Certificate(s): Any expired certificates on the switch.

Generate an SSL certificate

You can generate up to two self-signed SSL certificates for HTTPS communication.

> **NOTE:** Before you can generate a certificate, you must disable HTTPS (see Configure the HTTP access settings on page 581) and log back in to the main UI over an HTTP session. After you generate the certificate, you can reenable HTTPS and log back in to the main UI over an HTTPS session.

To generate an SSL certificate:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > HTTPS > Certificate Management**.
   The Certificate Management page displays.
3. From the **Certificate Number** menu, select the certificate, that is, select **1** or **2**.
   Your selection determines if the first or second certificate is generated.
4. Select the **Generate Certificates** radio button.
5. Click the **Apply** button.
   The switch generates an SSL certificate.
   The Certificate Generation Status field shows that the switch is in the process of generating a certificate.
6. Click the **Refresh** button.
   The Certificate Present field shows Yes, indicating that the certificate is present.
7. To save the settings to the running configuration, click the **Save** icon.

Activate a certificate

After you add a self-generated SSL certificate or transfer (upload) and external certificate authority (CA)-signed certificate to the switch, you can activate the certificate.

To activate a certificate:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > HTTPS > Certificate Management**.
   The Certificate Management page displays.
3. From the **Certificate Number** menu, select the certificate, that is, select **1** or **2**.
   Your selection determines if the first or second certificate is activated.
4. Select the **Activate Certificate** radio button.
5. Click the **Apply** button.
   Your settings are saved and the certificate now active.
6. To save the settings to the running configuration, click the **Save** icon.

Delete an SSL certificate

NOTE: Before you can delete a certificate, you must disable HTTPS (see Configure the HTTP access settings on page 581) and log back in to the main UI over an HTTP session. After you delete the certificate, you can reenable HTTPS and log back in to the main UI over an HTTPS session.

To delete an SSL certificate:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > HTTPS > Certificate Management**.
   The Certificate Present field displays Yes.
3. From the **Certificate Number** menu, select the certificate, that is, select **1** or **2**.
   Your selection determines if the first or second certificate is deleted.
4. Select **Delete Certificates** radio button.
5. Click the **Apply** button.
   The certificate is removed.

Transfer an existing HTTPS certificate from a server to the switch

You can transfer a certificate file from a secure server to the switch. A certificate that you transfer to the switch is assigned number 1.

NOTE: For information about downloading and installing an SSL certificate over an HTTP session, see Download and install an SSL security certificate file on the switch on page 779.

For the switch to accept HTTPS connections from a device, the switch requires a public key certificate. You can generate a certificate externally (for example, offline) or obtain a certificate authority (CA)-signed certificate and transfer it to the switch.
Before you transfer a file from a server to the switch, the following conditions must be true:

- The file that you transfer from a server is on the server in the appropriate directory.
- The file is in the correct format.
- The switch contains a path to the server.

NOTE: Before you can transfer a certificate, you must disable HTTPS (see Configure the HTTP access settings on page 581) and log back in to the main UI over an HTTP session. After you transfer the certificate, you can reenable HTTPS and log back in to the main UI over an HTTPS session.

To transfer an HTTPS certificate to the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > HTTPS > Certificate Download**.
   The Certificate Update page displays.
3. From the **File Type** menu, select the type of SSL certificate to download:
   - **SSL Trusted Root Certificate PEM File**: SSL Trusted Root Certificate file (PEM Encoded)
   - **SSL Server Certificate PEM File**: SSL Server Certificate File (PEM Encoded)
   - **SSL DH 1024 Encryption Parameter PEM File**: SSL Diffie-Hellman 1024-bit Encryption Parameter File (PEM Encoded)
   - **SSL DH 2048 Encryption Parameter PEM File**: SSL Diffie-Hellman 2048-bit Encryption Parameter File (PEM Encoded)
4. From the **Transfer Mode** menu, select the protocol that must be used to transfer the file:
   - **TFTP**: Trivial File Transfer Protocol
   - **SFTP**: Secure File Transfer Protocol
   - **SCP**: Secure Copy Protocol

5. From the **Server Address Type** menu, select **IPv4**, **IPv6,** or **DNS** to indicate the format for the TFTP Server IP field.
   The default is IPv4.
6. In the **Server Address** field, specify the address or host name of the server.
   The address can be an IP address in the standard IPv4 or IPv6 address format or a host name. The host name must start with a letter of the alphabet.
7. In the **Remote File Path** field, enter the path of the file to download.
   You can enter up to 96 characters. The default is blank.
8. In the **Remote File Name** field, enter the name of the file on the TFTP server to download.
   You can enter up to 32 characters. The default is blank.
9. If you select **SFTP** or **SCP** from the **Transfer Mode** menu, specify the following settings:
   - **User Name**: Specify the user name for login to the server.
   - **Password**: Specify the password for login to the server.
10. Click the **Apply** button.
    The file transfer starts.
    A status message displays during the transfer and upon successful completion of the transfer.
    After the file transfer is complete, the certificate is assigned number 1.
11. To save the settings to the running configuration, click the **Save** icon.

SSH management access

You can display and modify the Secure Shell (SSH) server settings on the switch. SSH is a network protocol that enables access to the CLI management interface by using an SSH client on a remote administrative system. SSH is a more secure access method than Telnet because it encrypts communication between the administrative system and the device. You can download or generate SSH host keys for secure CLI-based management.

Configure the global SSH access settings

You can configure the global SSH access settings.

To configure SSH settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > SSH > SSH Configuration**.
   The SSH Configuration page displays.

3. Select the SSH Admin Mode **Disable** or **Enable** radio button.
   This selection enables or disables the SSH global administrative mode. When enabled, you can access the switch by using an SSH client on a remote system. The default is Disable.
4. In the **SSH Session Timeout** field, specify the inactivity time-out period for SSH sessions to the switch.
   If an SSH session is inactive and the time-out period is exceeded, the session is terminated. The range is from 1 to 160 minutes. The default is 5 minutes.
5. In the **Maximum Number of SSH Sessions** field, specify the maximum number of SSH sessions that can be simultaneously made to the switch.
   The range is from 0 to 5. The default is 5 sessions.
6. From the **Authentication List** menu, select a login authentication list.
   This list is used to authenticate users who try to log in to the switch. The default list is networkList. For more information about login authentication lists, see Configure a login authentication list on page 574.
7. From the **Enable Authentication List** menu, select an enable authentication list.
   This list is used to authenticate users who try to get privileged EXEC access. The default list is enableList. For more information about enable authentication lists, see Configure an enable authentication list on page 576.
8. In the **SSH Port** field, specify the port for SSH access.
   The range is from 1 to 65535. The default port number is 22.
9. Select the SSH Public Key Authentication Mode **Disable** or **Enable** radio button.
   This selection enables or disables the SSH public key authentication mode. When enabled, you can use a secure public key rather than a password to access the switch over an SSH session. The default is Disable.
10. Select the SCP server Admin Mode **Disable** or **Enable** radio button.
    This selection enables or disables file transfers over a secure copy protocol (SCP) server. The default is Disable.
11. In the **Max SSH Authentication Retries** field, specify the number of times that the switch tries to authenticate an SSH client after the initial authentication fails.
    The range is from range from 0 to 5. The default is 3.
12. Click the **Apply** button.
    Your settings are saved.
13. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 175.SSH configuration information

Field	Description
SSH Version	The switch can accept connections from an SSH client using protocol level 2 for SSH (SSH-2). Protocol level 2 for SSH is enabled by default.
Current Number of SSH Sessions	The number of active SSH sessions between remote SSH clients and the switch.
Keys Present	Displays if one or both (if any) of the following keys are present on the device: • SSH-2 Rivest-Shamir-Adelman (RSA) key file (PEM encoded) • SSH-2 Digital Signature Algorithm (DSA) key file (PEM encoded) If no keys are presents, No is displayed.
Max SSH Authentication Tries	The number of times that the switch attempts to authenticate the SSH client. By default, this number is 4.

Manage RSA, DSA, and ECDSA keys for SSH access

You can manage Rivest-Shamir-Adelman (RSA) and various versions of Digital Signature Algorithm (DSA) keys for SSH access to the switch.

Generate an RSA, DSA, or ECDSA key

The switch supports the following keys that are used for SSH access:

- **RSA**: Rivest-Shamir-Adleman (RSA), which is an algorithm that is widely supported.
- **DSA**: Digital Signature Algorithm (DSA), which is a secure algorithm for authentication but not for encryption.
- **ECDSA**: Elliptic Curve Digital Signature Algorithm (ECDSA), which is a newer and very secure algorithm that is not yet widely supported.

NOTE: To generate an SSH key file, SSH must be disabled (see Configure the global SSH access settings on page 588).

To generate an RSA, DSA, or ECDSA key:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Access > SSH > Host Keys Management**.
   The Host Keys Management page displays.
   The Keys Present field displays whether a certificate is present on the switch.
3. Do one or more of the following:
   - **RSA key**: In the RSA Keys Management section, in the **Key Length** field, specify a key length, and select the **Generate RSA Keys** radio button. The supported key lengths for RSA are 1024, 2048 and 3072. The default is 1024.
   - **DSA key**: In the DSA Keys Management section, select the **Generate DSA Keys** radio button.
   - **ECDSA key**: In the ECDSA Keys Management section, in the **Key Length** field, specify a key length, and select the **Generate ECDSA Keys** radio button. The supported key lengths for ECDSA are 256, 384 and 521. The default is 256.
4. Click the **Apply** button.
   The switch generates a key. The Key Generation In Progress field shows progress information.
5. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 176.Host Key Status information

Field	Description
Keys Present	Displays which keys are present on the switch (if any): RSA, DSA, ECDSA, or None.
Key Generation In Progress	Displays which key is being generated (if any): RSA, DSA, ECDSA, or None.

Delete an RSA, DSA, or ECDSA key

NOTE: To delete an SSH key file, SSH must be disabled (see Configure the global SSH access settings on page 588).

To delete an RSA, DSA, or ECDSA key:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security** > **Access > SSH** > **Host Keys Management**.
   The Host Keys Management page displays.
   The Keys Present field displays Yes.
3. Do one or more of the following:
   - **RSA key**: In the RSA Keys Management section, select the **Delete RSA Keys** radio button.
   - **DSA key**: In the DSA Keys Management section, select the **Delete DSA Keys** radio button.
   - **ECDSA key**: In the ECDSA Keys Management section, select the **Delete ECDSA Keys** radio button.
4. Click the **Apply** button.
   The keys are removed.
5. To save the settings to the running configuration, click the **Save** icon.

Transfer existing SSH keys from a TFTP server to the switch

You can transfer an SSH-2 RSA, SSH-2 DSA , or ECDSA key Privacy Enhanced Mail (PEM) file from a remote switch or computer to the switch.
Before you transfer a file from a TFTP server to the switch, the following conditions must be true:

- The file that you transfer from a TFTP server is on the server in the appropriate directory.
- The file is in the correct format.
- The switch contains a path to the TFTP server.

NOTE: To download an SSH key file to the switch, SSH must be disabled (see Configure the global SSH access settings on page 588).

To transfer SSH keys to the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > SSH > Host Keys Download**.
   The Host Keys Download page displays.
3. From the **File Type** menu, select the type of key file to download:
   - **SSH-2 RSA Key PEM File**: SSH-2 Rivest-Shamir-Adelman (RSA) key file (PEM Encoded)
   - **SSH-2 DSA Key PEM File**: SSH-2 Digital Signature Algorithm (DSA) key file (PEM Encoded)
   - **ECDSA Key PEM File**: Elliptic Curve Digital Signature Algorithm (ECDSA) key file (PEM Encoded)

4. From the **Transfer Mode** menu, select the protocol that must be used to transfer the file:
   - **TFTP**: Trivial File Transfer Protocol
   - **SFTP**: Secure File Transfer Protocol
   - **SCP**: Secure Copy Protocol
5. From the **Server Address Type** menu, select **IPv4**, **IPv6,** or **DNS** to indicate the format for the TFTP Server IP field.
   The default is IPv4.
6. In the **Server IP** field, specify the address or host name of the server.
   The address can be an IP address in the standard IPv4 or IPv6 address format or a host name. The host name must start with a letter of the alphabet.
7. In the **Remote File Path** field, enter the path of the file to download.
   You can enter up to 96 characters. The default is blank.
8. In the **Remote File Name** field, enter the name of the file on the TFTP server to download.
   You can enter up to 32 characters. The default is blank.
9. If you select **SFTP** or **SCP** from the **Transfer Mode** menu, enter the user name in the **Username** field.
   This is the user name with which you access the SFTP or SCP server.
10. If you select **SFTP** or **SCP** from the **Transfer Mode** menu, enter the password in the **Password** field.
    This is the password that is required to access the SFTP or SCP server.
11. Click the **Apply** button.
    The file transfer starts. A status message displays during the transfer and upon successful completion of the transfer.
12. To save the settings to the running configuration, click the **Save** icon.

Telnet management access

You can configure Telnet authentication lists and manage outbound and inbound Telnet.

Select Telnet authentication lists

You can select a login authentication list and an enable authentication list:

- **Login authentication list**: The login list specifies the authentication methods used to validate switch or port access for the users associated with the list. For more information, see Configure a login authentication list on page 574.
- **Enable authentication list**: The enable list specifies the authentication methods used to validate privileged EXEC access for the users associated with the list. For more information, see Configure an enable authentication list on page 576.

To configure the Telnet authentication lists:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > Telnet**.
   The Telnet page displays.
3. From the **Authentication List** menu, select which login authentication list must be used to establish a Telnet session.
   The default is networkList.
4. From the **Enable Authentication List** menu, select which enable authentication list must be used to access the privileged EXEC mode in an established Telnet session.
   The default is enableList.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Configure inbound Telnet settings

If you enable inbound Telnet sessions and allow new Telnet sessions to be established, authorized users can establish new Telnet sessions to the switch until the maximum number of Telnet sessions is exceeded. An established Telnet session remains active until the session is ended or the session time-out period is exceeded.

To configure the inbound Telnet settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Access > Telnet**.
   The Telnet page displays.
3. In the Inbound Telnet section, configure the following settings:
   a. Select the Telnet Server Admin Mode **Disable** or **Enable** radio button.
      This selection specifies if the switch accepts *any* inbound Telnet sessions. By default, the **Enable** radio button is selected.
   b. Next to Allow new telnet sessions, select the **Disable** or **Enable** radio button.
      This selection specifies if the switch accepts *new* Telnet sessions, that is, new in addition to a current Telnet session. By default, the **Enable** radio button is selected. If you disable this option, an established session remains active until the session is ended, after which no new Telnet session can be established.
   c. In the **Session Timeout (Minutes)** field, specify the time in minutes after which an inactive Telnet session is automatically ended.
      The range is from 1 to 160 minutes. The default is 5 minutes.
   d. In the **Maximum Number of Sessions** field, specify how many simultaneous inbound Telnet sessions are allowed.
      The maximum is 5, which is also the default.
4. Click the **Apply** button.
   Your settings are saved.
   The Current Number of Sessions field in the Inbound Telnet section displays the number of current inbound Telnet sessions.
5. To save the settings to the running configuration, click the **Save** icon.

Configure outbound Telnet settings

If you enable outbound Telnet sessions and allow new Telnet sessions to be established, authorized users can establish new Telnet sessions from the switch to another device until the maximum number of Telnet sessions is exceeded. An established Telnet session remains active until the session is ended or the session time-out period is exceeded.

To configure the outbound Telnet settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > Telnet**.
   The Telnet page displays.

3. In the Outbound Telnet section, configure the following settings:
   a. Next to Allow new telnet sessions, select the **Disable** or **Enable** radio button.
      This selection specifies if the switch allows *new* outbound Telnet sessions. By default, the **Enable** radio button is selected. If you disable this option, an established session remains active until the session is ended.
   b. In the **Session Timeout (Minutes)** field, specify the time in minutes after which an inactive Telnet session is automatically ended.
      The range is from 1 to 160 minutes. The default is 5 minutes.
   c. In the **Maximum Number of Sessions** field, specify how many simultaneous outbound Telnet sessions are allowed.
      The maximum is 5, which is also the default.
4. Click the **Apply** button.
   Your settings are saved.
   The Current Number of Sessions field in the Outbound Telnet section displays the number of current outbound Telnet sessions.
5. To save the settings to the running configuration, click the **Save** icon.

Console port management access

You can configure the settings for a serial port (console port) connection to the switch.

To configure the console port settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > Console Port**.
   The Console Port page displays.
3. In the **Serial Port Login Timeout (minutes)** field, specify the time in minutes after which an inactive console port connection is automatically ended.
   The range is from 1 to 160 minutes. The default is 5 minutes. Entering 0 disables the time-out.
4. From the **Baud Rate (bps)** menu, select the default baud rate for the console port connection.
   You can choose from 1200, 2400, 4800, 9600, 19200, 38400, 57600, and 115200 baud. The default is 115200 baud.

5. From the **Authentication List** menu, select a login authentication list.
   This list is used to authenticate users who try to log in to the switch. The default is defaultList. For more information about login authentication lists, see Configure a login authentication list on page 574.
6. From the **Enable Authentication List** menu, select an enable authentication list.
   This list is used to authenticate users who try to get privileged EXEC access. The default list is enableList. For more information about enable authentication lists, see Configure an enable authentication list on page 576.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 177.Console port information

Field	Description
Character Size (bits)	The number of bits in a character. This setting is always 8.
Flow Control	Shows if hardware flow control is enabled or disabled. This setting is always disabled.
Stop Bits	The number of stop bits per character. This setting is always 1.
Parity	The parity method used on the console port. This setting is always None.

Denial of service

You can select which types of denial of service (DoS) attacks the switch monitors and blocks.

To configure individual DoS settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > Denial of ServiceConfiguration**.
   The Denial of Service Configuration page displays.

3. Select the types of DoS attacks for the switch to monitor and block and configure any associated values:
   - **Denial of Service Min TCP Header Size**: Specify the minimum TCP header size allowed. If you select the **Denial of Service TCP Fragment** radio button, the switch drops the first TCP fragment with a TCP payload packet for which the minimum TCP header size is larger than the IP payload length minus the IP header size. The range for the minimum TCP header size is from 0 to 255. The default value is 20.
   - **Denial of Service ICMPv4**: Enabling ICMPv4 DoS prevention causes the switch to drop ICMPv4 packets with a type set to ECHO_REQ (ping) and a size greater than the configured ICMPv4 packet size. By default, this option is disabled.
   - **Denial of Service Max ICMPv4 Packet Size**: Specify the maximum ICMPv4 packet size allowed. If ICMPv4 DoS prevention is enabled, the switch drops ICMPv4 ping packets with a size greater than the configured value. The range is from 0 to 16376. The default value is 512.
   - **Denial of Service ICMPv6**: Enabling ICMPv6 DoS prevention causes the switch to drop ICMPv6 packets with a type set to ECHO_REQ (ping) and a size greater than the configured ICMPv6 packet size. By default, this option is disabled.
   - **Denial of Service Max ICMPv6 Packet Size**: Specify the maximum ICMPv6 packet size allowed. If ICMPv6 DoS prevention is enabled, the switch drops ICMPv6 ping packets with a size greater than the configured value. The range is from 0 to 16376. The default value is 512.
   - **Denial of Service First Fragment**: Enabling first fragment DoS prevention causes the switch to check the DoS options on the first fragment of a fragmented IP packet. Otherwise, the switch ignores the first fragment of a fragmented IP packet. By default, this option is disabled.
   - **Denial of Service ICMP Fragment**: Enabling ICMP Fragment DoS prevention causes the switch to drop ICMP fragmented packets. By default, this option is disabled.
   - **Denial of Service SIP=DIP**: Enabling SIP=DIP DoS prevention causes the switch to drop packets with a source IP address equal to the destination IP address. By default, this option is disabled.
   - **Denial of Service SMAC=DMAC**: Enabling SMAC=DMAC DoS prevention causes the switch to drop packets with a source MAC address equal to the destination MAC address. By default, this option is disabled.
   - **Denial of Service TCP FIN&URG&PSH**: Enabling TCP FIN & URG & PSH DoS prevention causes the switch to drop packets with TCP flags FIN, URG, and PSH set and the TCP sequence number equal to 0. By default, this option is disabled.
   - **Denial of Service TCP Flag&Sequence**: Enabling TCP Flag DoS prevention causes the switch to drop packets with TCP control flags set to 0 and the TCP sequence number set to 0. By default, this option is disabled.
   - **Denial of Service TCP Fragment**: Enabling TCP Fragment DoS prevention causes the

switch to drop packets with a TCP payload for which the IP payload length minus the IP header size is less than the minimum allowed TCP header size. By default, this option is disabled.

- **Denial of Service TCP Offset**: Enabling TCP Offset DoS prevention causes the switch to drop packets with a TCP header offset set to 1. By default, this option is disabled.
- **Denial of Service TCP Port**: Enabling TCP Port DoS prevention causes the switch to drop packets for which the TCP source port is equal to the TCP destination port. By default, this option is disabled.
- **Denial of Service TCP SYN**: Enabling TCP SYN DoS prevention causes the switch to drop packets with the TCP flag SYN set. By default, this option is disabled.
- **Denial of Service TCP SYN&FIN**: Enabling TCP SYN & FIN DoS prevention causes the switch to drop packets with TCP flags SYN and FIN set. By default, this option is disabled.
- **Denial of Service UDP Port**: Enabling UDP Port DoS prevention causes the switch to drop packets for which the UDP source port is equal to the UDP destination port. By default, this option is disabled.
- **Stacked VLAN0 tag drop mode**: Enabling the tag drop mode for stacked VLAN 0 causes the switch to forward or drop packets that are tagged with VLAN 0 and forward or drop SNAP packets. (SNAP stands for Stateful Network-Wide Abstractions for Packet Processing.) By default, this option is disabled.

4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Management access profiles and rules

Access control allows you to configure an access control profile and set rules for access to the main UI, access by SNMP stations, and client access to a TFTP server. We refer to an access control profile as an access profile. You can add a single access profile, which you can configure, activate, or deactivate.

⚠ **CAUTION:** If you configure a security access profile incorrectly and you activate the access profile, you might no longer be able to access the switch's main UI. If that situation occurs, you must reset the switch to factory default settings (see Reset the switch to the factory default settings on page 767.)

Add an access profile

You can set up a single security access profile with which you can associate an access rule configuration.

To add an access profile:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > Access Control > Access Profile Configuration**.
   The Access Profile Configuration page displays.
3. In the **Access Profile Name** field, enter the name of the access profile to be added.
   The maximum length is 32 characters.
4. Click the **Apply** button.
   Your settings are saved. By default, the access profile is deactivated. After you add rules, you can activate the access profile.
5. To save the settings to the running configuration, click the **Save** icon.

Add a rule to the access profile

After you add the access profile, you can add one or more security access rules to the access profile.
If you access the switch from a computer, make sure that you add a permit rule for the type of service that you use (for example, HTTPS), your computer's IP address, and your computer's subnet mask.

⚠ **CAUTION:** You must add a permit rule for your device and access method, otherwise you are locked out from the switch after you activate the access profile. If that situation occurs, you must reset the switch to factory default settings (see Reset the switch to the factory default settings on page 767).

To add a rule to the access profile:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > Access Control > Access Rule Configuration**.
   The Access Rule Configuration page displays.

3. From the **Rule Type** menu, select **Permit** or **Deny** to permit or deny access when the selected rules are matched.
   A Permit rule allows access from a device that matches the rule criteria. A Deny rule blocks a device that matches the rule criteria.
4. From the **Service Type** menu, select the access method to which the rule is applied.
   Unless you allow any access method, the policy is restricted by the selected access method. Possible access methods are **Any**, **TELNET**, **TFTP**, **HTTP,Secure HTTP (SSL)**, **SNMP**, **Secure Telnet (SSH)**, and **NTP**.
5. In the **Source IP Address** field, enter the source IP address from which the management traffic originates.
6. In the **Mask** field, specify the subnet mask from which the management traffic originates.
7. In the **Priority** field, assign a priority to the rule.
   The rules are validated against the incoming management request in ascending order of their priorities. If a rule matches, the action is performed and subsequent rules below that rule are ignored. For example, if a source IP address 10.10.10.10 is configured with priority 1 to permit, and the same source IP address 10.10.10.10 is also configured with priority 2 to deny, then access is permitted if the profile is active, and the second rule is ignored.
8. Click the **Add** button.
   The access rule is added.
9. To save the settings to the running configuration, click the **Save** icon.

Activate the access profile

After you add rules to the access profile, you can activate the access profile.



CAUTION: If you configure a security access profile incorrectly and you activate the access profile, you might no longer be able to access the switch's main UI. If that situation occurs, you must reset the switch to factory default settings (see Reset the switch to the factory default settings on page 767).

To activate the access profile:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > Access Control > Access Profile Configuration**.
   The Access Profile Configuration page displays. The Deactivate Profile check box is selected.

3. Select the **Activate Profile** check box.
4. Click the **Apply** button.
   Your settings are saved and the access profile is now active.
5. To save the settings to the running configuration, click the **Save** icon.

Display the access profile summary and the number of filtered packets

After you added rules to the active profile, you can display the entries in the summary. If the access profile is active, you can also display the number of filtered packets.

To display the access profile summary and the number of filtered packets:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > Access Control > Access Profile Configuration**.
   The Access Profile Configuration page displays.
   The Packets Filtered field displays the number of packets filtered.
3. To refresh the page with the latest information about the switch, click the **Refresh** button.
   The following table describes the view-only fields on the page..

Table 178.Access profile configuration profile summary

Field	Description
Rule Type	The action performed when the rules match.
Service Type	The service type selected. The policy is restricted by the selected service type.
Source IP Address	The source IP address of the client originating the management traffic.
Mask	The subnet mask of the IP Address.
Priority	The priority of the rule.

Deactivate an access profile

You can deactivate an access profile.

To deactivate an access profile:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > Access Control > Access Profile Configuration**.
   The Access Profile Configuration page displays. The Activate Profile check box is selected.
3. Select the **Deactivate Profile** check box.
4. Click the **Apply** button.
   Your settings are saved and the access profile is now deactivated.
5. To save the settings to the running configuration, click the **Save** icon.

Remove an access profile

You can remove an access profile that you no longer need. Before you can remove the access profile, you must deactivate it (see Deactivate an access profile on page 602).

To remove an access profile:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > Access Control > Access Profile Configuration**.
   The Access Profile Configuration page displays. The Deactivate Profile check box is selected.
3. Select the **Remove Profile** check box.
4. Click the **Apply** button.
   The access profile is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Port authentication

With port-based authentication, when 802.1X is enabled both globally and on the port, successful authentication of any one supplicant attached to the port results in all users being able to use the port without restrictions. At any time, only one supplicant is allowed to attempt authentication on a port in this mode. Ports in this mode are under bidirectional control. 802.1X is the default authentication mode. 802.1X is also referred to as dot1x.

NOTE: For port authentication, if we refer to a port, it means the same as a physical interface.

An 802.1X network includes three components:

- **Authenticator**: The port that is authenticated before access to system services is permitted.
- **Supplicant**: The host that is connected to the authenticated port requesting access to the system services.
- **Authentication server**: The external server, for example, the RADIUS server that performs the authentication on behalf of the authenticator, and indicates whether the supplicant is authorized to access system services.

Configure the global 802.1X authentication settings

You can enable 802.1X on the switch and configure the global 802.1X settings that apply to the switch and to specific users.
If you enable 802.1X, authentication must be performed by a RADIUS server:

- **RADIUS server**: Configure a RADIUS server (see RADIUS servers on page 562).
- **Primary authentication method**: Set the primary authentication method to RADIUS, that is, RADIUS must be method 1 for the defaultList setting of the login authentication list (see Configure a login authentication list on page 574).

To configure the global 802.1X settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Port Authentication > Basic > 802.1X Configuration**.
   The 802.1X Configuration page displays.
3. Select the Dot1x Admin Mode **Disable** or **Enable** radio button:
   - **Enabled**: 802.1X port-based authentication is globally enabled on the switch.
   - **Disabled**: 802.1X port-based authentication is globally disabled on the switch. This is the default setting. The switch does not check for 802.1X authentication before allowing traffic on any ports, even if individual ports are configured to allow only authenticated users.

4. Select the Authentication Admin Mode **Disable** or **Enable** radio button:
   - **Enabled**: The authentication admin mode is enabled, requiring authentication to be performed by a RADIUS server. This means that the primary authentication method must be RADIUS. To set the method, select Security > Management Security > Authentication List > Authentication List and select **Radius** as method 1 for defaultList. For more information, see Configure a login authentication list on page 574.
   - **Disabled**: The authentication admin mode is disabled, which is the default setting. The switch does not check for 802.1X authentication before allowing traffic on any ports, even if the 802.1X port-based authentication is globally enabled (see the previous step) and the ports are configured to allow only authenticated users.
5. Select the VLAN Assignment Mode **Disable** or **Enable** radio button.
   When enabled, this feature allows a port to be placed into a particular VLAN based on the result of the authentication or type of 802.1X authentication a client uses when it accesses the device. The authentication server can provide information to the device about which VLAN to assign the supplicant. The default is Disable.
6. Select the EAPOL Flood Mode **Disable** or **Enable** radio button.
   This selection specifies whether Extensible Authentication Protocol (EAP) over LAN (EAPoL) flood support is enabled on the switch. The default is Disable.
7. From the **Dynamic VLAN Creation Mode** menu, select **Disable** or **Enable**.
   If RADIUS-assigned VLANs are enabled, the RADIUS server includes the VLAN ID in the 802.1X tunnel attributes of its response message to the device. If dynamic VLAN creation is enabled on the device and the RADIUS-assigned VLAN does not exist, the assigned VLAN is dynamically created. This means that the client can connect from any port and is assigned to the appropriate VLAN. This feature gives flexibility for clients to move around the network without much additional configuration required. The default is Disable.
8. From the **Monitor Mode** menu, select **Disable** or **Enable**.
   If enabled, this option lets the switch monitor the dot1x authentication process and can help you to diagnose authentication failures, if they occur. The default is Disable.
9. From the **Users** menu, select the user name to which the login list that you can select in the next step must apply.
10. From the **Login** menu, select the login list that applies to the user name that you select in the previous step.
    If 802.1X is enabled, the default dot1xList applies for authentication, as displayed in the Authentication List field.
    The Software Version field always display 0.
11. Click the **Apply** button.
    Your settings are saved.

12. To save the settings to the running configuration, click the **Save** icon.

Manage port authentication on individual ports

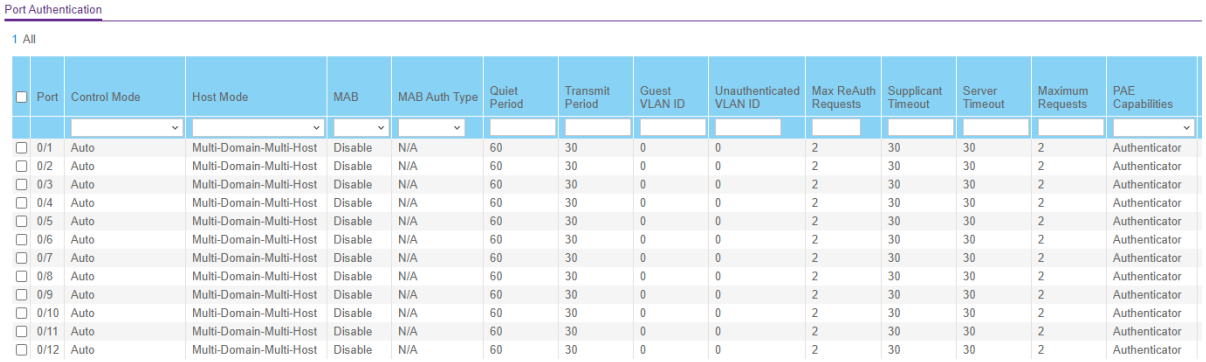
You can enable and configure port access control on one or more physical ports.

Configure 802.1X settings for a port

You can configure 802.1X port access control settings for one or more ports.

To configure 802.1X settings for a port:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Port Authentication > Advanced > Port Authentication**.


Port Authentication

1 All

Port	Control Mode	Host Mode	MAB	MAB Auth Type	Quiet Period	Transmit Period	Guest VLAN ID	Unauthenticated VLAN ID	Max ReAuth Requests	Supplicant Timeout	Server Timeout	Maximum Requests	PAE Capabilities
0/1	Auto	Multi-Domain-Multi-Host	Disable	N/A	60	30	0	0	2	30	30	2	Authenticator
0/2	Auto	Multi-Domain-Multi-Host	Disable	N/A	60	30	0	0	2	30	30	2	Authenticator
0/3	Auto	Multi-Domain-Multi-Host	Disable	N/A	60	30	0	0	2	30	30	2	Authenticator
0/4	Auto	Multi-Domain-Multi-Host	Disable	N/A	60	30	0	0	2	30	30	2	Authenticator
0/5	Auto	Multi-Domain-Multi-Host	Disable	N/A	60	30	0	0	2	30	30	2	Authenticator
0/6	Auto	Multi-Domain-Multi-Host	Disable	N/A	60	30	0	0	2	30	30	2	Authenticator
0/7	Auto	Multi-Domain-Multi-Host	Disable	N/A	60	30	0	0	2	30	30	2	Authenticator
0/8	Auto	Multi-Domain-Multi-Host	Disable	N/A	60	30	0	0	2	30	30	2	Authenticator
0/9	Auto	Multi-Domain-Multi-Host	Disable	N/A	60	30	0	0	2	30	30	2	Authenticator
0/10	Auto	Multi-Domain-Multi-Host	Disable	N/A	60	30	0	0	2	30	30	2	Authenticator
0/11	Auto	Multi-Domain-Multi-Host	Disable	N/A	60	30	0	0	2	30	30	2	Authenticator
0/12	Auto	Multi-Domain-Multi-Host	Disable	N/A	60	30	0	0	2	30	30	2	Authenticator

   The previous figure shows part of the page only.

3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

5. From the **Control Mode** menu, select an option:
   - **Auto**: The authenticator port access entity (PAE) sets the controlled port mode to reflect the outcome of the authentication exchanges between the supplicant, authenticator, and the authentication server. This is the default setting.
   - **Force Authorized**: The authenticator PAE unconditionally sets the controlled port to authorized.
   - **Force Unauthorized**: The authenticator PAE unconditionally sets the controlled port to unauthorized.

6. From the **Host Mode** menu, select an option to specify the number of clients and the type of clients that can be authenticated and authorized on the port:
   - **Single-Host**: One data client only can be authenticated on the port. After authentication succeeds, access is granted to this client only but not to other clients. Only when this client logs off can another client be authenticated and authorized on the port and granted access to the port.
   - **Multi-Host**: Initially, one data client only can be authenticated on the port. After authentication succeeds, access is granted to all clients connected to the port. As an example, use this option when a WiFi access point is connected to an access-controlled port of a NAS. After the access point is authenticated by the NAS, the port is authorized for traffic from not just the access point but also from all the WiFi clients connected to the access point.
   - **Multi-Domain**: One data client and one voice client can be authenticated on the port. After authentication succeeds, the data and voice clients are granted access. As an example, use this option when an IP phone is connected to a NAS port and a laptop is connected to the hub port of the IP phone. Both devices must be authenticated to access the network services behind the NAS. The voice and data domains are segregated. (The RADIUS server attribute Cisco-AVPair = device-traffic-class=voice is used to identify a voice client.)
   - **Multi-Auth**: One voice client and multiple data clients can be authenticated on the port. After authentication succeeds, access is granted to all clients. As an example, use this option when a network of laptops and an IP phone are connected to a NAS port via a hub.
   - **Multi-Domain-Multi-Host**: Initially, one voice client and one data client can be authenticated on the port. After the data client is authenticated, access is granted to all clients connected to the port and they are considered data clients. As an example, use this option when an IP phone is connected to a NAS port and a virtual machine (VM) controller is connected to the hub port of the IP phone. The VM controller hosts multiple VMs. Both the VM controller and the IP phone must be authenticated to access the network services behind the NAS. The voice and data domains are segregated. After the VM controller is authenticated, traffic is allowed from all VMs hosted by the VM controller. Note that if the data client is authenticated first, the voice client can be authenticated only using 802.1x.
7. From the **MAB** menu, select to enable or disable MAC-based authentication bypass (MAB) for 802.1x unaware clients.
   MAB functions only if the port control mode is MAC-based. The default selection is Disable.

8. From the **MAB Auth Type** menu, select a MAB authentication option:
   - **EAP-MD5**: The MD5 hash of the MAC address is sent as the password in the EAP message (Radius Attribute 79) to the authentication server.
   - **PAP**: The MAC address of the client is sent as the password, similar to the format of Attribute 1, in clear text as part of the User-Password message (Radius Attribute 2).
   - **CHAP**: A randomly generated 16-octet challenge is sent as the CHAP-Challenge message (Radius Attribute 60) along with the CHAP-Password message (Radius Attribute 3). The CHAP ID is a unique number that is used to identify the session. The MAC address of the client is retrieved and formatted using the configured Attribute 1 format. Then, this information is used as a secret to derive the information for the CHAP-Password message. The information for the CHAP-Password message is calculated as MD5 (with the CHAP-ID, secret, and CHAP-Challenge).
9. In the **Quiet Period** field, enter the period in seconds during which the interface does not attempt to acquire a supplicant after an earlier authentication exchange failed.
   Enter a value in the range from 0 to 65535. A quiet period of 0 means that the interface does not acquire a supplicant at all. The default is 60 seconds.
10. In the **Transmit Period** field, enter the period in seconds after which the interface sends an EAPOL EAP Request/Identity frame to the supplicant.
    Enter a value in the range from 0 to 65535. The default is 30 seconds.
11. In the **GuestVLAN ID** field, enter the ID of the guest VLAN.
    Enter a value in the range from 0 to 4093. The default is 0.
12. In the **Unauthenticated VLAN ID** field, enter the ID for the unauthenticated VLAN ID
    A user is allowed three attempts to enter the correct credentials. Otherwise, the client is placed in the unauthenticated VLAN. Enter a value in the range from 0 to 4093. The default is 0.
13. In the **Max ReAuth Requests** field, enter the maximum number of reauthentication requests that are allowed.
    Enter a value in the range from 1 to 20. The default value is 2.
14. In the **Supplicant Timeout** field, enter the period in second after which the interface times out the supplicant.
    Enter a value in the range from 0 to 65535. The default is 30 seconds.
15. In the **Server Timeout** field, enter the period after which the interface times out the authentication server.
    Enter a value in the range from 0 to 65535. The default is 30 seconds.

16. In the **Maximum Requests** field, enter the maximum number of times that the interface sends an EAPOL EAP request/identity to the supplicant before timing out the supplicant.
    Enter a value in the range from 1 to 10. The default is 2.
17. From the **PAE Capabilities** menu, select **Authenticator** or **Supplicant** to specify the function of the port access entity (PAE).
    The default is Authenticator.
18. From the **Periodic Reauthentication** menu, select to enable or disable the periodic reauthentication option of the supplicant.
    The default is Disable.
19. If you enable the periodic reauthentication option, in the **Reauthentication Period** field, enter the period in seconds after which the supplicant must be reauthenticated.
    Enter a value in the range from 0 to 65535. The default is 3600 seconds.
20. From the **User Privileges** menu, select **admin** or **guest** to limit the type of users that can be granted access.
    By default, both admin and guest users can be granted access.
21. In the **Max Users** field, enter the maximum number of supplicants that the interface can support.
    The default is 48.
22. In the Authentication Order columns, enter the order in which authentication must occur for up to three methods:
    - **Method 1**: The default selection is DOT1X. The other options are MAB, and CAPTIVE PORTAL. You cannot disable Method 1.
    - **Method 2**: The default selection is MAB. The other options are None, DOT1X, and CAPTIVE PORTAL. Disable Method 2 by selecting **None**.
    - **Method 3**: The default selection is CAPTIVE PORTAL. The other options are None, DOT1X, and MAB. Disable Method 3 by selecting **None**.

    As an example, if a client tries 802.1X authentication (based on the user and client credentials) and authentication fails, the interface proceeds with MAB authentication. If that method fails too, the interface proceeds with captive portal authentication. If that method fails too, the client cannot be authenticated.
23. Click the **Apply** button.
    Your settings are saved.
24. To save the settings to the running configuration, click the **Save** icon.

Initialize 802.1X on a port

If you initialize 802.1X on a port, the port becomes unauthorized and then goes through the authentication procedure, causing traffic on the port to be blocked until the port is authorized successfully.
You can initialize 802.1X on a port only if the control mode setting is Auto. This is the default setting, in which the authenticator port access entity (PAE) sets the controlled port mode to reflect the outcome of the authentication exchanges between the supplicant, authenticator, and the authentication server.

To initialize 802.1X on a port:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Port Authentication > Advanced > Port Authentication**.
   The Port Authentication page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. Click the **Initialize** button.
   802.1X on the selected port is reset to the initialization state. Traffic sent to and from the port is blocked during the authentication process. This button is available only if the control mode is auto. When you click this button, the action is immediate. You do not need to click the **Apply** button for the action to occur.

Display the port summary

You can display summary information about the port-based authentication settings for each port.

To display the port summary:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Port Authentication > Advanced > Port Summary**.
   The Port Summary page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 179.Port summary information

Field	Description
Port	The port for which settings are displayed
Control Mode	The configured control mode for the port: ○ **Force Unauthorized**: The authenticator port access entity (PAE) unconditionally sets the controlled port to unauthorized. ○ **Force Authorized**: The authenticator PAE unconditionally sets the controlled port to authorized. ○ **Auto**: The authenticator PAE sets the controlled port mode to reflect the outcome of the authentication exchanges between the supplicant, authenticator, and the authentication server.
Operating Control Mode	The control mode under which the port is actually operating. The options are as follows: ○ ForceUnauthorized ○ ForceAuthorized ○ Auto ○ MAC Based ○ N/A: If the port is in detached state, the value is N/A because the port cannot participate in port access control.
Reauthentication Enabled	Indicates if reauthentication of the supplicant is allowed. If the value is True, reauthentication is allowed. Otherwise, reauthentication is not allowed.
Control Direction	The control direction, which determines the degree to which protocol exchanges occur between a supplicant and authenticator. The direction affects whether the unauthorized controlled port exerts control over communication in both directions (disabling both incoming and outgoing frames) or just in the incoming direction (disabling only the reception of incoming frames).
PAE Capabilities	The port access entity (PAE) functionality, which is either Authenticator or Supplicant
Authenticator PAE State	The state of the authenticator PAE: ○ Initialize ○ Authenticating ○ Authenticated ○ Held ○ Unauthenticated

Field	Description
Key Transmission Enabled	Indicates if key transmission is enabled. If the value is False, key transmission does not occur. Otherwise, key transmission is supported.
Port Status	The authorization status of the port: ○ N/A: If the port is in detached state, the value is N/A because the port cannot participate in port access control. ○ Authorized ○ Unauthorized
Protocol Version	The protocol version is always 2, corresponding to the second version of the 802.1X specification.

Display the client summary

You can display information about supplicant devices that are connected to the local authenticator ports. If no active 802.1X sessions exist, the table is empty.

To display the client summary:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Port Authentication > Advanced > Client Summary**.
   The Client Summary page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   ○ **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   ○ **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.

4. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 180.Client summary information

Field	Description
Port	The port for which information is displayed
User Name	The name that represents the identity of the supplicant device
Supplicant Mac Address	The MAC address of the supplicant's device
Session Time	The period in seconds since the supplicant logged in
Filter ID	The policy filter ID assigned by the authenticator to the supplicant device
VLAN ID	The VLAN ID assigned by the authenticator to the supplicant device
VLAN Assigned	The reason for the VLAN ID assigned by the authenticator to the supplicant device: ○ **Default Assigned VLAN**: The client was authenticated on the port default VLAN and the authentication server was not a RADIUS server ○ **Radius Assigned VLAN**: RADIUS was used to authenticate the client ○ **Unauthenticated VLAN:** The client was authenticated on the Unauthenticated VLAN ○ **Guest VLAN**: The client was authenticated on the Guest VLAN ○ **Voice VLAN**: The client was authenticated on the Voice VLAN ○ **Monitor Mode VLAN**: The client was authenticated in Monitor mode and assigned by the RADIUS server to a monitor VLAN ○ **Not Assigned**: The client was not assigned to any VLAN
Session Timeout	The session time-out enforced by the RADIUS server for the supplicant device
Termination Action	The termination action enforced by the RADIUS server for the supplicant device
Time left for Session Termination Action	The time left before the session is terminated for the reason that is displayed in the Termination Action field

MAC filters for traffic control

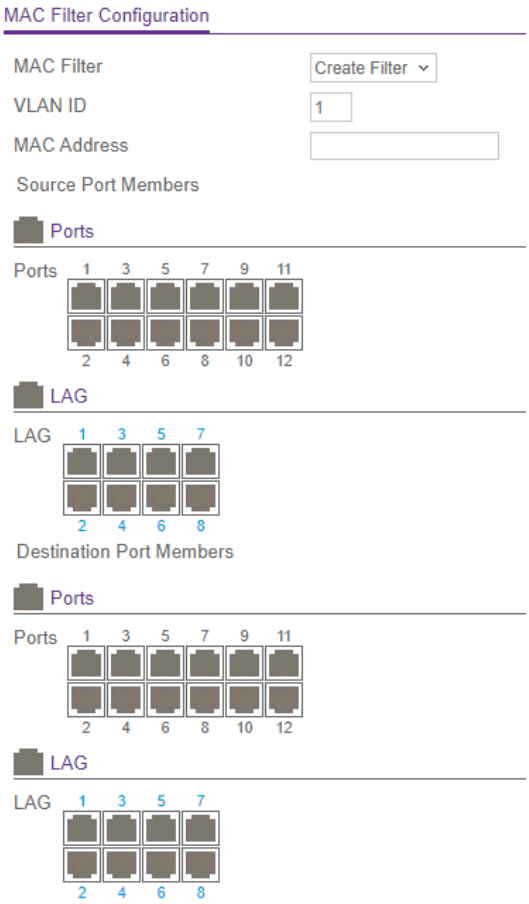
You can create MAC filters that limit the traffic allowed into and out of specified ports on the switch. The traffic limitations are based on MAC addresses.

Create a MAC filter

You can create MAC filters that limit the traffic allowed into and out of specified ports on the system.

To create a MAC filter:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > MAC Filter**.



3. From the **MAC Filter** menu, select **Create Filter**.
   If you did not configure any filters, this is the only option available.
4. From the **VLAN ID** menu, select the VLAN that must be used with the MAC address.

5. In the **MAC Address** field, specify the MAC address of the filter in the format XX:XX:XX:XX:XX:XX.
   You cannot define filters for the following MAC addresses:
   - 00:00:00:00:00:00
   - 01:80:C2:00:00:00 to 01:80:C2:00:00:0F
   - 01:80:C2:00:00:20 to 01:80:C2:00:00:21
   - FF:FF:FF:FF:FF:FF
6. In the Ports table (or if a stack is configured, in the Ports tables) and in LAG table in the Source Port Members section, select the ports and LAGs that must be included in the inbound filter.
   The ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed, and you can select ports on different stacked switches.
   If a packet with the MAC address and VLAN ID that you specify is received on a port or LAG that is not part of the inbound filter, the packet is dropped.
7. In the Ports table (or if a stack is configured, in the Ports tables) and in the and LAG table in the Destination Port Members section, select the ports and LAGs that must be included in the outbound filter.
   Also in this section, the ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed, and you can select ports on different stacked switches.
   A packet with the MAC address and VLAN ID that you specify can be transmitted only from a port or LAG that is part of the outbound filter.

   NOTE: Destination ports can be included only in a multicast filter. A multicast filter is determined by the MAC address that you enter in the MAC Address field.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Delete a MAC filter

You can remove an existing MAC filter that you no longer need.

To delete a MAC filter:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Traffic Control > MAC Filter**.
   The MAC Filter Configuration page displays.
3. From the **MAC Filter** menu, select the MAC filter.
4. Click the **Delete** button.
   The filter is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Display the MAC filter summary

You can display the MAC filters that are configured on the switch.

To display the MAC filter summary:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > MAC Filter > MAC Filter Summary**.
   The MAC Filter Summary page displays.
   The following table describes the view-only fields on the page.

Table 181.MAC filter summary information

Field	Description
MAC Address	The MAC address of the filter in the format XX:XX:XX:XX:XX:XX.
VLAN ID	The VLAN ID used with the MAC address to fully identify packets you want filtered.
Source Port Members	The ports to be used for filtering inbound packets.
Destination Port Members	The ports to be used for filtering outbound packets.

Port security

Port security lets you lock one or more ports on the switch. When a port is locked, the port can only forward packets if the total number of dynamically learned MAC addresses and the total number of statically added MAC addresses are not exceeded. You can set the thresholds for these numbers. Once a threshold is exceeded, the port discards all other packets.
You can also convert dynamically learned MAC addresses to static MAC addresses and allow traffic from these MAC addresses on a port.

Configure the global port security mode

Before you can enable and configure port security for individual ports, you must globally enable the port security mode for the switch.

To configure the global port security mode:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Port Security > Port Administration**.
   The Port Administration page displays.
3. To enable port security on the switch, select the Port Security Mode **Enable** radio button.
   The default is Disable.
4. Click the **Apply** button.
   Your settings are saved.
   The Port Security Violations table shows information about violations that occurred on ports that are enabled for port security.

Table 182.Port security violations information

Field	Description
Port	The port for which the violation is displayed.
Last Violation MAC	The source MAC address of the last packet that was discarded on the locked port.
VLAN ID	The VLAN ID corresponding to the last MAC address violation.

Configure a port security interface

A MAC address can be defined as allowable by one of two methods: dynamically or statically. Both methods are used concurrently when a port is locked.
Dynamic locking implements a first arrival mechanism for port security. You specify how many addresses can be learned on the locked port. If the limit is not reached, a packet with an unknown source MAC address is learned and forwarded normally. If the limit is reached, no more addresses are learned on the port. Any packets with source MAC addresses that are not already learned are discarded. You can effectively disable dynamic locking by setting the number of allowable dynamic entries to zero.
Static locking allows you to specify a list of MAC addresses that are allowed on a port. The behavior of packets is the same as for dynamic locking: only packets with an allowable source MAC address can be forwarded.

To configure port security settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Port Security > Interface Configuration**.
   The Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

5. Specify the following settings:
   - **Security Mode**: Enable or disable the port security feature. The default is Disable.
   - **Max Learned MAC Address**: Specify the maximum number of dynamically learned MAC addresses. The default value is 4096. If you specify 0, the interface does not learn any MAC addresses.
   - **Max Static MAC Address**: Specify the maximum number of statically locked MAC addresses. The default value is 48.
   - **Violation Shutdown**: Enable or disable shutdown if a packet with a disallowed MAC address is received. The default is Disable.
   - **Violation Trap**: Enable or disable sending of violation traps if a packet with a disallowed MAC address is received. The default is Disable.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Display learned MAC addresses and convert them to static addresses

After you enable port security globally (see Configure the global port security mode on page 619) and enable port security for specific interfaces (see Configure a port security interface on page 620), you can convert a dynamically learned MAC address to a statically locked address.

To display learned MAC addresses for an individual interface or LAG and convert these MAC addresses to static MAC addresses:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Port Security > Dynamic MAC Address**.
   The Dynamic MAC Address page displays.

3. To display the learned MAC addresses for an interface, from the **Port List** menu, select the interface or LAG.
   The Number of Dynamic MAC Addresses Learned field displays the number of dynamically learned MAC addresses on a specific port.
   The following table shows the MAC addresses and their associated VLANs learned on the selected interface.

Field	Description
VLAN ID	The VLAN ID corresponding to the MAC address.
MAC Address	The MAC addresses learned on port.

4. To convert the dynamically learned MAC addresses on the interface to statically locked addresses, select the **Convert Dynamic Address to Static** check box.
   The dynamic MAC address entries are converted to static MAC address entries in a numerically ascending order until the static limit is reached.
5. To refresh the page, click the **Refresh** button.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Add a static MAC address to the MAC address table for port security

Static MAC address entries are the ones that you manually add to the MAC address table for port security for a specific interface and VLAN.

To add a static MAC address to the MAC address table for port security:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Port Security > Static MAC Address**.
   The Static MAC Address page displays.
3. From the **Interface** menu, select the interface or LAG.
4. In the **Static MAC Address** field, enter the MAC address.
5. From the **VLAN ID** menu, select the VLAN ID that must be associated with the MAC address.
6. Click the **Add** button.
   The static MAC address is added to the MAC address table.

7. To save the settings to the running configuration, click the **Save** icon.

Remove a static MAC address from the MAC address table for port security

You can remove a static MAC address that you no longer need for port security.

To remove a static MAC address from the MAC address table for port security:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Port Security > Static MAC Address**.
   The Static MAC Address page displays.
3. From the **Interface** menu, select the interface or LAG.
4. Select the check box for the MAC address.
5. Click the **Delete** button.
   The static MAC address is removed from the MAC address table.
6. To save the settings to the running configuration, click the **Save** icon.

Private port groups

To add a level of security, you can set up a private group of physical ports (not LAGs) that can function in isolated or community mode:

- **Isolated mode group**: A port that you add as member of the group cannot forward its egress traffic to any other members in the same group.
- **Community mode group**: A port that you add as member of the group *can* forward its egress traffic to other members in the same group, but not to members in other groups.

Add a private port group

After you add a private group, you can add members to the group (see Configure the membership of a private port group on page 624).

To add a private port group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Traffic Control > Private Group > Private Group Configuration**.
   The Private Group Configuration page displays.
3. In the **Group Name** field, enter the private group name.
   The name can be up to 24 characters.
4. From the **Group Mode** menu, select the mode for the private group:
   - **Isolated**: A port that you add as member of the group cannot forward its egress traffic to any other members in the same group.
   - **Community**: A port that you add as member of the group *can* forward its egress traffic to other members in the same group, but not to members in other groups.
5. Click the **Add** button.
   The private group is added.
6. To save the settings to the running configuration, click the **Save** icon.

Remove a private port group

You can remove a private port group that you no longer need.

To remove a private port group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Private Group > Private Group Configuration**.
   The Private Group Configuration page displays.
3. Select the check box next to the private group.
4. Click the **Delete** button.
   The private group is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the membership of a private port group

After you add a private port group, you can add members to the group (see Add a private port group on page 623).

To configure the membership of a private port group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Traffic Control > Private Group > Private Group Membership**.
   The Private Group Membership page displays.
3. From the **Group ID** menu, select the group ID.
4. In the Ports table (or if a stack is configured, in one of the Ports tables), select the ports that must be members of the group.
   The ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed, and you can select ports on different stacked switches.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 183.Private group membership information

Field	Description
Group Name	The name of the private group
Group Mode	The mode of the private group (Community or Isolated)

Protect ports

If a port is configured as protected, it does not forward traffic to any other protected port on the switch, but it does forward traffic to unprotected ports. You can configure ports as protected or unprotected.

To configure protected ports:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Protected Ports**.
   The Protected Ports page displays.
3. From the **Group ID** list, select **0**, **1**, or **2** to combine the selected protected ports in a logical group.
   Traffic can flow between protected ports belonging to different groups, but not within the same group. The group ID can be 0, 1, or 2.

4. To associate a name with the protected ports group, specify a name in the **Group Name** field.
   The name can be up to 32 alphanumeric characters long, including blanks.
5. In the Ports table (or if a stack is configured, in one of the Ports tables), select the ports that must be protected ports and members of the protected port group.
   The ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed, and you can select ports on different stacked switches.
   By default, all ports are unprotected.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Private VLANs

A private VLAN contains switch ports that cannot communicate with each other, but can access another network. These ports are called private ports. Each private VLAN contains one or more private ports and a single uplink port or uplink aggregation group. Note that all traffic between private ports is blocked at all layers, not just Layer 2 traffic, but also traffic such as FTP, HTTP, and Telnet.
A private VLAN separates a regular VLAN domain into two or more subdomains. Each subdomain is defined (represented) by a primary VLAN and a secondary VLAN:

- **Primary VLAN**: The primary VLAN ID is the same for all subdomains that belong to a private VLAN.
- **Secondary VLAN**: The secondary VLAN ID differentiates subdomains from each other and provides Layer 2 isolation between ports of the same private VLAN.

Within a private VLAN, three types of VLANs can exist:

- **Primary VLAN**: The VLAN forwards traffic from promiscuous ports to isolated ports, community ports, and other promiscuous ports in the same private VLAN. In a private VLAN, you can configure one primary VLAN only. All ports in a private VLAN share the same primary VLAN.
- **Isolated VLAN**: The VLAN is a secondary VLAN that carries traffic from isolated ports to promiscuous ports. In a private VLAN, you can configure one isolated VLAN only.
- **Community VLAN**: The VLAN is a secondary VLAN that forwards traffic between ports that belong to the same community and to the promiscuous ports. In a private VLAN, you can configure multiple community VLANs.

Within a private VLAN, the switch supports four types of special port designations:

- **Host port**: The port is a host port that is a member of a community VLAN or an isolated VLAN, both of which are secondary VLANs within the private VLAN. Two host port subtypes

exist:

- **Community port**: The port is a member of a community VLAN. A community port can communicate with other community ports and promiscuous ports.
- **Isolated port**: The port is a member of an isolated VLAN. An isolated port can communicate with promiscuous ports.

- **Promiscuous port**: The port is a member of a primary VLAN (within the private VLAN) and can communicate with all types of ports in the private VLAN, including other promiscuous ports, community ports, and isolated ports.
- **Promiscuous trunk port**: The port is both an uplink trunk port (for example, it can connect a router) and a promiscuous port that can be a member of multiple primary VLANs (within the private VLAN) as well as regular VLANs. The port processes tagged traffic to communicate with all types of ports in private VLANs (including promiscuous ports, community ports, and isolated ports) and with ports in regular VLANs.
- **Isolated trunk port**: The port is both a downlink trunk port (for example, it can connect to multiple downstream devices) and can be a member of multiple isolated VLANs (secondary VLANs within private VLANs) as well as regular VLANs. The port processes tagged traffic to communicate with ports in isolated VLANs and with ports in regular VLANs. The downstream devices that are connected to the isolated trunk port do not need to be capable of supported private VLANs.

Overview of the tasks for private VLAN configuration

To set up a private VLAN that allows for communication between switches in a network, perform the tasks that are described in the following sections:

1. Assign a private VLAN type to a VLAN on page 628.
   By default, a VLAN is a regular VLAN, so you must assign a private VLAN type to a VLAN.
2. Configure a private VLAN association with a primary and secondary VLAN on page 629.
   A private VLAN must consist of a single primary VLAN and one or more secondary VLANs.
3. Configure the private VLAN port mode on page 631.
   A port must be in the correct port mode to participate in a private VLAN. For example, to make a port a member of an isolated VLAN or community VLAN, you must first configure the port as a host port.
4. Private VLAN host interface: Assign the interface to primary and secondary VLANs on page 633.
   For a port that you configured to function in host mode, configure a single primary VLAN and a single secondary VLAN.

5. Private VLAN promiscuous interface: Assign the interface to primary and secondary VLANs on page 635.
   For a port that you configured to function in promiscuous mode, configure a single primary VLAN and one or more secondary VLANs.
6. Private VLAN promiscuous trunk interface: Add primary and secondary VLANs to the trunk on page 637.
   For a port that you configured to function in promiscuous *trunk* mode, add primary and secondary VLANs to the trunk.
7. Private VLAN isolated trunk interface: Add primary and secondary VLANs to the trunk on page 639
   For a port that you configured to function in isolated *trunk* mode, add primary and secondary VLANs to the trunk.
8. Configure native and allowed VLANs on a private VLAN trunk interface on page 641
   For port that you configured to function in a private VLAN *trunk* mode (either in promiscuous trunk mode or in isolated trunk mode), configure native VLANs for untagged traffic and allowed private VLANs.

Assign a private VLAN type to a VLAN

To each VLAN, you can assign a private VLAN type, which can be Primary, Isolated, or Community. By default, a VLAN is a regular VLAN and assigned the private VLAN type Unconfigured.

To assign a private VLAN type to a VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Private VLAN > Private VLAN Type Configuration**.
   The Private VLAN Type Configuration page displays.
3. Select the check box for the VLAN.

4. From the **Private VLAN Type** menu, select the type of private VLAN:
   - **Unconfigured**: The VLAN is not a private VLAN but a regular VLAN. This the default setting.
   - **Primary**: The VLAN is a primary VLAN that forwards traffic from promiscuous ports to isolated ports, community ports, and other promiscuous ports in the same private VLAN. In a private VLAN, you can configure one primary VLAN only. All ports in a private VLAN share the same primary VLAN.
   - **Isolated**: The VLAN is a secondary VLAN that carries traffic from isolated ports to promiscuous ports. In a private VLAN, you can configure one isolated VLAN only.
   - **Community**: The VLAN is a secondary VLAN that forwards traffic between ports that belong to the same community and to the promiscuous ports. In a private VLAN, you can configure multiple community VLANs.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Configure a private VLAN association with a primary and secondary VLAN

You can configure a private VLAN by associating a single primary VLAN with one or more secondary VLANs.

To configure a private VLAN association with a primary and secondary VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Private VLAN > Private VLAN Association Configuration**.
   The Private VLAN Association Configuration page displays.
3. If you are changing an existing private VLAN association, select the check box for the private VLAN association.
4. From the **Primary VLAN** menu, select the primary VLAN ID for the private VLAN.
   This selection specifies the primary VLAN within the private VLAN. You can associate secondary VLANs in the private VLAN with this primary VLAN.
   For information about configuring a primary VLAN for the private VLAN, see Assign a private VLAN type to a VLAN on page 628.

5. In the **Secondary VLANs** field, specify one or more secondary VLANs for the private VLAN.
   This selection specifies secondary VLANs (isolated VLANs, community VLANs, or a combination of both) within the private VLAN. The secondary VLANs are associated with the primary VLAN in the private VLAN.
   You can specify a single VLAN ID, a range of VLAN IDs, or a combination of both in sequence separated by a comma (,):
   - You can specify individual VLAN ID, such as 10.
   - You can specify the VLAN range values separated by a hyphen, for example, 10-13.
   - You can specify the combination of both separated by commas, for example: 12,15,40–43,1000–1005, 2000.

   For information about configuring an isolated or community VLAN for the private VLAN, see Assign a private VLAN type to a VLAN on page 628.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 184.Private VLAN Association

Field	Description
Isolated VLAN	The single isolated VLAN associated with the selected primary VLAN.
Community VLAN(s)	The list of community VLANs associated with the selected primary VLAN.

Remove an existing private VLAN association

You can remove a private VLAN association that you no longer need.

To remove a private VLAN association:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Private VLAN > Private VLAN Association Configuration**.
   The Private VLAN Association Configuration page displays.
3. Select the check box for the private VLAN association.

4. Click the **Delete** button.
   Your settings are saved. The private VLAN association is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the private VLAN port mode

A port must be in the correct port mode to participate in a private VLAN.
The private plan port mode determines if a port (or LAG) can function in Host mode for a primary or secondary VLAN (within a private VLAN) or in Promiscuous mode for a promiscuous VLAN (within a private VLAN).
You can also configure a port (or LAG) to function as a promiscuous trunk for an uplink to a router or as an isolated trunk for a downlink to multiple devices. These types of private VLAN trunk ports can carry multiple private VLANs as well as regular VLANs.

To configure the private VLAN port mode:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Private VLAN > Private VLAN Port Mode Configuration**.
   The Private VLAN Port Mode Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.

5. From the **Port VLAN Mode** menu, select the private VLAN port mode:
   - **General**: The port functions in general mode and not as a port in a private VLAN. This is the default setting.
   - **Host**: The port functions in host mode in a private VLAN. In this mode, the port can be member of a community VLAN or isolated VLAN:
     - **Community VLAN**: The port is a member of a secondary VLAN (within the private VLAN) and can communicate with other community ports and promiscuous ports.
     - **Isolated VLAN**: The port is a member of a secondary VLAN (within the private VLAN) and can communicate only with promiscuous ports.

     To configure a host port to be a member of specific VLANs, see Private VLAN host interface: Assign the interface to primary and secondary VLANs on page 633.
   - **Promiscuous**: The port functions in promiscuous mode in a private VLAN. In this mode, the port can communicate with all types of ports in the private VLAN, including other promiscuous ports, community ports, and isolated ports. To configure a promiscuous port to be a member of specific VLANs, see Private VLAN promiscuous interface: Assign the interface to primary and secondary VLANs on page 635.
   - **Promiscuous trunk port**: The port functions in promiscuous trunk mode as both an uplink trunk port (for example, it can connect a router) and a promiscuous port that can be a member of multiple primary VLANs (within the private VLAN) as well as regular VLANs. In this mode, the port processes tagged traffic to communicate with all types of ports in private VLANs (including promiscuous ports, community ports, and isolated ports) and with ports in regular VLANs. To assign a promiscuous trunk port to specific VLANs, see Private VLAN promiscuous trunk interface: Add primary and secondary VLANs to the trunk on page 637.
   - **Isolated trunk port**: The port functions in isolated trunk mode (for example, it can connect to multiple downstream devices) and can be a member of multiple isolated VLANs (secondary VLANs within private VLANs) as well as regular VLANs. In this mode, the port processes tagged traffic to communicate with ports in isolated VLANs and with ports in regular VLANs. The downstream devices that are connected to the isolated trunk port do not need to be capable of supported private VLANs. To assign an isolated trunk port to specific VLANs, see Private VLAN isolated trunk interface: Add primary and secondary VLANs to the trunk on page 639.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Private VLAN host interface: Assign the interface to primary and secondary VLANs

If you configure the private VLAN port mode of an interface as Host (see Configure the private VLAN port mode on page 631), you can assign the interface to a single primary VLAN and single secondary VLAN.

To assign a private VLAN host interface to a primary and secondary VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Private VLAN > Private VLAN Host Interface Configuration**.
   The Private VLAN Host Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the interface, or type the interface number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
5. In the **Host Primary VLAN** field, specify a primary VLAN ID.
   You can select a VLAN for which you configured the type as Primary (see Assign a private VLAN type to a VLAN on page 628).
6. In the **Host Secondary VLAN** field, specify a secondary VLAN ID.
   You can select a VLAN for which you configured the type as Isolated or Community, both of which are secondary VLAN types within a private VLAN (see Assign a private VLAN type to a VLAN on page 628).

7. Click the **Apply** button.
   Your settings are saved.
   The Operational VLAN(s) fields shows the primary and secondary VLANs that operate on the host interface.
8. To save the settings to the running configuration, click the **Save** icon.

Private VLAN host interface: Remove the interface from primary and secondary VLANs

You can remove a private VLAN host interface from primary and secondary VLANs.

To remove a private VLAN host interface from primary and secondary VLANs:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Private VLAN > Private VLAN Host Interface Configuration**.
   The Private VLAN Host Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.

   The interface that you select must be configured in Host mode (see Configure the private VLAN port mode on page 631).
5. Click the **Delete** button.
   Your settings are saved. The interface is removed from the primary and secondary VLANs.
6. To save the settings to the running configuration, click the **Save** icon.

Private VLAN promiscuous interface: Assign the interface to primary and secondary VLANs

If you configure the private VLAN port mode of an interface as Promiscuous (see Configure the private VLAN port mode on page 631), you can assign the interface to a single primary VLAN and to one or more secondary VLANs.

To assign a private VLAN promiscuous interface to a primary VLAN and secondary VLANs:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Private VLAN > Private VLAN Promiscuous Interface Configuration**.
   The Private VLAN Promiscuous Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.

   The interface that you select must be configured in Promiscuous mode (see Configure the private VLAN port mode on page 631).
5. In the **Promiscuous Primary VLAN** field, specify a primary VLAN ID.
   You can select a VLAN for which you configured the type as Primary (see Assign a private VLAN type to a VLAN on page 628).

6. In the **Promiscuous Secondary VLAN IDs** field, specify one or more secondary VLAN IDs.
   You can specify VLANs for which you configured the type as Isolated or Community, both of which are secondary VLAN types within a private VLAN (see Assign a private VLAN type to a VLAN on page 628).
   You can specify a single VLAN ID, a range of VLAN IDs, or a combination of both in sequence separated by a comma (,):
   - You can specify individual VLAN ID, such as 10.
   - You can specify the VLAN range values separated by a hyphen, for example, 10-13.
   - You can specify the combination of both separated by commas, for example: 12,15,40–43,1000–1005, 2000.

   NOTE: The VLAN IDs that you specify overwrites the secondary VLANs that you configured on the Private VLAN Association Configuration page (see Configure a private VLAN association with a primary and secondary VLAN on page 629).

7. Click the **Apply** button.
   Your settings are saved.
   The Operational VLAN(s) fields shows the primary and secondary VLANs that operate on the promiscuous interface.

8. To save the settings to the running configuration, click the **Save** icon.

Private VLAN promiscuous interface: Remove the interface from primary and secondary VLANs

You can remove a private VLAN promiscuous interface from primary and secondary VLANs.

To remove a private VLAN promiscuous interface from primary and secondary VLANs:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Traffic Control > Private VLAN > Private VLAN Promiscuous Interface Configuration**.
   The Private VLAN Promiscuous Interface Configuration page displays.

3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.

   The interface that you select must be configured in Promiscuous mode (see Configure the private VLAN port mode on page 631).
5. Click the **Delete** button.
   Your settings are saved.
   Your settings are saved. The interface is removed from the primary and secondary VLANs.
6. To save the settings to the running configuration, click the **Save** icon.

Private VLAN promiscuous trunk interface: Add primary and secondary VLANs to the trunk

If you configure the private VLAN port mode of an interface as Promiscuous Trunk (see Configure the private VLAN port mode on page 631), you can add multiple VLAN configurations to the trunk. Each VLAN configuration consists of single primary VLAN and one or more secondary VLANs.

To add a primary VLAN and secondary VLANs to a private VLAN promiscuous trunk interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Private VLAN > Private VLAN Promiscuous Trunk Interface Configuration**.
   The Private VLAN Promiscuous Trunk Interface Configuration page displays.

3. From the **Interface** menu, select an interface or LAG for which you configured the VLAN port mode as Promiscuous Trunk mode.
   For more information about port VLAN modes, see Configure the private VLAN port mode on page 631.
4. In the **Promiscuous Trunk Primary VLAN** field, specify a primary VLAN ID.
   You can select a VLAN for which you configured the type as Primary (see Assign a private VLAN type to a VLAN on page 628).
5. In the **Promiscuous Trunk Secondary VLAN(s)** field, specify one or more secondary VLAN IDs.
   You can specify VLANs for which you configured the type as Isolated or Community, both of which are secondary VLAN types within a private VLAN (see Assign a private VLAN type to a VLAN on page 628).
   You can specify a single VLAN ID, a range of VLAN IDs, or a combination of both in sequence separated by a comma (,):
   - You can specify individual VLAN ID, such as 10.
   - You can specify the VLAN range values separated by a hyphen, for example, 10-13.
   - You can specify the combination of both separated by commas, for example: 12,15,40–43,1000–1005, 2000.

   NOTE: The VLAN IDs that you specify overwrites the secondary VLANs that you configured on the Private VLAN Association Configuration page (see Configure a private VLAN association with a primary and secondary VLAN on page 629).
6. Click the **Apply** button.
   Your settings are saved.
   The Operational VLAN(s) fields shows the primary and secondary VLANs that operate on the promiscuous trunk interface.
7. To save the settings to the running configuration, click the **Save** icon.

Private VLAN promiscuous trunk interface: Remove primary and secondary VLANs from the trunk

You can remove primary and secondary VLANs from a private VLAN promiscuous trunk interface.

To remove primary and secondary VLANs from a private VLAN promiscuous trunk interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Traffic Control > Private VLAN > Private VLAN Promiscuous Trunk Interface Configuration**.
   The Private VLAN Promiscuous Trunk Interface Configuration page displays.
3. From the **Interface** menu, select an interface or LAG for which you configured the VLAN port mode as Promiscuous Trunk mode.
   For more information about port VLAN modes, see Configure the private VLAN port mode on page 631.
4. Select the check box for the VLAN configuration.
5. Click the **Delete** button.
   Your settings are saved.
   The primary and secondary VLANs are removed from the selected interface. The VLANs themselves are not deleted.
6. To save the settings to the running configuration, click the **Save** icon.

Private VLAN isolated trunk interface: Add primary and secondary VLANs to the trunk

If you configure the private VLAN port mode of an interface as Isolated Trunk (see Configure the private VLAN port mode on page 631), you can add multiple VLAN configurations to the trunk. Each VLAN configuration consists of single primary VLAN and one or more secondary VLANs.

To add a primary VLAN and secondary VLANs to a private VLAN isolated trunk interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Private VLAN > Private VLAN Isolated Trunk Interface Configuration**.
   The Private VLAN Isolated Trunk Interface Configuration page displays.
3. From the **Interface** menu, select an interface or LAG for which you configured the VLAN port mode as Isolated Trunk mode.
   For more information about port VLAN modes, see Configure the private VLAN port mode on page 631.
4. In the **Isolated Trunk Primary VLAN** field, specify a primary VLAN ID.
   You can select a VLAN for which you configured the type as Primary (see Assign a private VLAN type to a VLAN on page 628).

5. In the **Isolated Trunk Secondary VLAN(s)** field, specify one or more secondary VLAN IDs.
   You can specify VLANs for which you configured the type as Isolated or Community, both of which are secondary VLAN types within a private VLAN (see Assign a private VLAN type to a VLAN on page 628).
   You can specify a single VLAN ID, a range of VLAN IDs, or a combination of both in sequence separated by a comma (,):
   - You can specify individual VLAN ID, such as 10.
   - You can specify the VLAN range values separated by a hyphen, for example, 10-13.
   - You can specify the combination of both separated by commas, for example: 12,15,40–43,1000–1005, 2000.

   NOTE: The VLAN IDs that you specify overwrites the secondary VLANs that you configured on the Private VLAN Association Configuration page (see Configure a private VLAN association with a primary and secondary VLAN on page 629).

6. Click the **Apply** button.
   Your settings are saved.
   The Operational VLAN(s) fields shows the primary and secondary VLANs that operate on the isolated trunk interface.

7. To save the settings to the running configuration, click the **Save** icon.

Private VLAN isolated trunk interface: Remove primary and secondary VLANs from the trunk

You can remove primary and secondary VLANs from a private VLAN isolated trunk interface.

To remove primary and secondary VLANs from a private VLAN isolated trunk interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Traffic Control > Private VLAN > Private VLAN Isolated Trunk Interface Configuration**.
   The Private VLAN Isolated Trunk Interface Configuration page displays.

3. From the **Interface** menu, select an interface or LAG for which you configured the VLAN port mode as Isolated Trunk mode.
   For more information about port VLAN modes, see Configure the private VLAN port mode on page 631.

4. Select the check box for the VLAN configuration.

5. Click the **Delete** button.
   Your settings are saved.
   The primary and secondary VLANs are removed from the selected interface. The VLANs themselves are not deleted.
6. To save the settings to the running configuration, click the **Save** icon.

Configure native and allowed VLANs on a private VLAN trunk interface

If you configure an interface as either a private VLAN promiscuous interface or a private VLAN isolated trunk interface, you can also allow untagged traffic from a native VLAN on the same trunk interface. By default, VLAN 1 is both the default and the native VLAN, but you can configure another VLAN as the native VLAN (see Configure the switch port mode settings for interfaces on page 201).
If you do not specify a native VLAN on the private trunk interface, all untagged packets are dropped from the private VLAN trunk interface.

To configure native and allowed VLANs on a private VLAN trunk interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Private VLAN > Private VLAN Trunk Interface Configuration**.
   The Private VLAN Trunk Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.

4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.

   The interface that you select must be configured either in Promiscuous Trunk mode or Isolated Trunk mode (see Configure the private VLAN port mode on page 631).
5. In the **Trunk Native VLAN** field, specify a VLAN ID for a native VLAN.
   By default, VLAN 1 is both the default and the native VLAN, but you can configure another VLAN as the native VLAN (see Configure the switch port mode settings for interfaces on page 201).
6. In the **Trunk Allowed VLAN** field, specify one or more VLAN IDs for private VLAN configurations that consists of primary VLAN and secondary VLANs.
   You can specify VLANs for which you configured the type as Promiscuous, Isolated, or Community within a private VLAN (see Assign a private VLAN type to a VLAN on page 628).
   You can specify a single VLAN ID, a range of VLAN IDs, or a combination of both in sequence separated by a comma (,):
   - You can specify individual VLAN ID, such as 10.
   - You can specify the VLAN range values separated by a hyphen, for example, 10-13.
   - You can specify the combination of both separated by commas, for example: 12,15,40–43,1000–1005, 2000.

   NOTE: The VLAN IDs that you specify overwrites the secondary VLANs that you configured on the Private VLAN Association Configuration page (see Configure a private VLAN association with a primary and secondary VLAN on page 629).
7. Click the **Apply** button.
   Your settings are saved.
   The Operational VLAN(s) field shows the VLANs that operate on the private VLAN trunk interface.
8. To save the settings to the running configuration, click the **Save** icon.

Storm control

A broadcast storm is the result of an excessive number of broadcast messages simultaneously transmitted across a network by a single port. Forwarded message responses can overload network resources, cause the network to time out, or do both.
The switch measures the incoming packet rate per port for broadcast, multicast, unknown, and unicast packets and discards packets if the rate exceeds the defined value. You enable storm

control per interface, by defining the packet type and the rate at which the packets are transmitted.

Configure global storm control settings

The global storm control settings apply to all ports. After you configure the global settings, you can specify storm control settings for one or more ports.

To configure global storm control settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Storm Control > Storm Control Global Configuration**.
   The Storm Control Global Configuration page displays.
   The following three settings allow you to globally enable or disable each type of packet to be rate-limited on all ports. For information about individual ports, see Configure storm control settings for one or more ports on page 644.
3. Configure the storm control settings by selecting a radio button:
   - **Broadcast Storm Control All**: Enable of disable this option:
     - **Enable**: If broadcast traffic on a port exceeds the configured threshold, the switch discards the broadcast traffic. The default is Enable.
     - **Disable**: Broadcast storm control is disabled on all ports.
   - **Multicast Storm Control All**: Enable of disable this option:
     - **Enable**: If multicast traffic on a port exceeds the configured threshold, the switch discards the multicast traffic.
     - **Disable**: Multicast storm control is disabled on all ports. The default is Disable.
   - **Unknown Unicast Storm Control All**: Enable of disable this option:
     - **Enable**: If unknown unicast traffic on a port exceeds the configured threshold, the switch discards the unknown unicast traffic.
     - **Disable**: Unknown unicast storm control is disabled on all ports. The default is Disable.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Configure storm control settings for one or more ports

You can specify storm control settings for one or more ports.

To configure storm control settings for one or more ports:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Traffic Control > Storm Control > Storm Control Interface Configuration**.
   The Storm Control Interface Configuration page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

5. Configure the following settings for broadcast storms:
   - **Recovery Mode**: Enable of disable this option for broadcast traffic:
     - **Enable**: If broadcast traffic exceeds the configured threshold, the port discards the broadcast traffic. The default is Enable.
     - **Disable**: If broadcast traffic exceeds the configured threshold, broadcast is not discarded.
   - **Recovery Level Type**: Specify the recovery level in the **Recovery Level** field as a percentage of the port link speed or as packets per second:
     - **Percent**: The recovery level is expressed as a percentage of link speed.
     - **pps**: The recovery level is expressed as a packets per second (pps).
   - **Recovery Level**: Specify the threshold at which broadcast storm control is activated. By default, the recovery level is 5 percent of the port link speed of 5 pps.
   - **Control Action**: Specify one of the following actions:
     - **None**: No further action is taken. This is the default setting.
     - **Trap**: If broadcast traffic exceeds the configured threshold, a trap is sent.
     - **Shutdown**: If broadcast traffic exceeds the configured threshold, the port is shut down.
6. Configure the following settings for multicast storms:
   - **Recovery Mode**: Enable of disable this option for multicast traffic:
     - **Enable**: If multicast traffic exceeds the configured threshold, the port discards the multicast traffic. The default is Enable.
     - **Disable**: If multicast traffic exceeds the configured threshold, multicast is not discarded.
   - **Recovery Level Type**: Specify the recovery level in the **Recovery Level** field as a percentage of the port link speed or as packets per second:
     - **Percent**: The recovery level is expressed as a percentage of link speed.
     - **pps**: The recovery level is expressed as a packets per second (pps).
   - **Recovery Level**: Specify the threshold at which multicast storm control is activated. By default, the recovery level is 5 percent of the port link speed of 5 pps.
   - **Control Action**: Specify one of the following actions:
     - **None**: No further action is taken. This is the default setting.
     - **Trap**: If multicast traffic exceeds the configured threshold, a trap is sent.
     - **Shutdown**: If multicast traffic exceeds the configured threshold, the port is shut down.

7. Configure the following settings for unicast storms:
   - **Recovery Mode**: Enable of disable this option for unicast traffic:
     - **Enable**: If unicast traffic exceeds the configured threshold, the port discards the unicast traffic. The default is Enable.
     - **Disable**: If unicast traffic exceeds the configured threshold, unicast is not discarded.
   - **Recovery Level Type**: Specify the recovery level in the **Recovery Level** field as a percentage of the port link speed or as packets per second:
     - **Percent**: The recovery level is expressed as a percentage of link speed.
     - **pps**: The recovery level is expressed as a packets per second (pps).
   - **Recovery Level**: Specify the threshold at which unicast storm control is activated. By default, the recovery level is 5 percent of the port link speed of 5 pps.
   - **Control Action**: Specify one of the following actions:
     - **None**: No further action is taken. This is the default setting.
     - **Trap**: If unicast traffic exceeds the configured threshold, a trap is sent.
     - **Shutdown**: If unicast traffic exceeds the configured threshold, the port is shut down.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

DHCP snooping

DHCP snooping is a feature that provides security by filtering untrusted DHCP messages and by building and maintaining a DHCP snooping binding table. An untrusted message is a message that is received from outside the network or firewall and that can cause traffic attacks within your network. The DHCP snooping binding table contains the MAC address, IP address, lease time, binding type, VLAN number, and interface information that corresponds to the local untrusted interfaces of a switch. An untrusted interface is an interface that is configured to receive messages from outside the network or firewall. A trusted interface is an interface that is configured to receive only messages from within the network.
DHCP snooping acts like a firewall between untrusted hosts and DHCP servers. It also provides a way to differentiate between untrusted interfaces connected to the end user and trusted interfaces connected to the DHCP server or another switch.

Enable DHCP snooping for the switch

You can globally enable DHCP snooping for the switch.

To globally enable DHCP snooping for the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCP Snooping > Global Configuration**.
   The Global Configuration page displays.
3. Select the DHCP Snooping Mode **Enable** radio button.
   The default is Disable.
4. To enable the verification of the sender's MAC address for DHCP snooping, leave the MAC Address Validation **Enable** radio button selected.
   The default is Enable.
   When MAC address validation is enabled, the switch checks packets that are received on an untrusted interface to verify that the MAC address and the DHCP client hardware address match. If the addresses do not match, the switch drops the packet.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Enable DHCP snooping for a VLAN

You can enable DHCP snooping on one or more VLANs.

To enable DHCP snooping for a VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCP Snooping > Global Configuration**.
   The Global Configuration page displays.
3. In the **VLAN ID** field, type a VLAN ID.
4. From the **DHCP Snooping Mode** menu, select **Enable**.
5. Click the **Apply** button.
   Your settings are saved.
   The VLAN ID displays in the VLAN Configuration table.
6. To save the settings to the running configuration, click the **Save** icon.

Configure DHCP snooping interface settings

You can display and configure each port as a trusted or untrusted port. Any DHCP responses received on a trusted port are forwarded. If a port is configured as untrusted, any DHCP (or BootP) responses received on that port are discarded.

To configure DHCP snooping interface settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCP Snooping > Interface Configuration**.
   The Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.

5. From the **Trust Mode** menu, select the trust mode:
   - **Disabled**: The interface is considered to be untrusted and could potentially be used to launch a network attack. DHCP server messages are checked against the bindings database. On untrusted ports, DHCP snooping enforces the following security rules:
     - DHCP packets from a DHCP server are dropped.
     - DHCP messages are dropped if the MAC address is in the snooping database but the binding's interface is other than the interface where the message was received.
     - DHCP packets are dropped if the source MAC address does not match the client hardware address and if MAC address validation is globally enabled.
   - **Enabled**: The interface is considered to be trusted and forwards DHCP server messages without validation.
6. From the **Invalid Packets** menu, select the packet logging mode.
   When enabled, the DHCP snooping feature generates a log message when an invalid packet is received and dropped by the interface.
7. In the **Rate Limit (pps)** field, specify the rate limit value for DHCP snooping purposes.
   If the incoming rate of DHCP packets per second exceeds the configured burst interval per second, the port shuts down. If the rate limit is None (which is the default), the burst interval is also not applicable, and rate limiting is disabled.
8. In the **Burst Interval (secs)** field, specify the burst interval in seconds for rate limiting on the interface.
   If the rate limit is N/A, the burst interval is not applicable.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

Add a static DHCP binding and display or clear dynamic DHCP bindings

You can add a static binding in the DHCP snooping bindings database and display or clear the dynamic bindings in the bindings table.

To add a static DHCP binding and display or clear the dynamic bindings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCP Snooping > Binding Configuration**.
   The Binding Configuration page displays.

3. From the **Interface** menu, select the interface.
4. In the **MAC Address** field, specify the MAC address for the binding to be added.
   This is the key to the binding database.
5. From the **VLAN ID** menu, select the ID of the VLAN.
6. In the **IP Address** field, specify the IP address for the binding to be added.
7. Click the **Add** button.
   The DHCP snooping binding entry is added to the database.
   The Dynamic Binding Configuration table shows information about the DHCP bindings that were learned on each interface on which DHCP snooping is enabled. The following table describes the dynamic binding information.

Field	Description
Interface	The interface on which the DHCP client message was received.
MAC Address	The MAC address associated with the DHCP client that sent the message. This is the key to the binding database.
VLAN ID	The VLAN ID of the client interface.
IP Address	The IP address assigned to the client by the DHCP server.
Lease Time	The remaining IP address lease time for the client.

8. To save the settings to the running configuration, click the **Save** icon.
9. To clear all dynamic bindings from the Dynamic Binding Configuration table, click the **Clear** button.
10. To save the settings to the running configuration, click the **Save** icon.

Remove a static DHCP binding

You can remove a static binding from the DHCP snooping bindings database.

To remove a static DHCP binding:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCP Snooping > Binding Configuration**.
   The Binding Configuration page displays.
3. In the Static Binding Configuration table, select the check box for the static binding.

4. Click the **Delete** button.
   The DHCP snooping binding entry is removed from the database.
5. To save the settings to the running configuration, click the **Save** icon.

Configure DHCP snooping persistent settings

You can configure the persistent location of the DHCP snooping bindings database. The bindings database can be stored locally on the switch or on a remote device in the network. The switch must be able to reach the IP address of the remote device to send bindings to a remote database.

To configure DHCP snooping persistent settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCP Snooping > Persistent Configuration**.
   The Persistent Configuration page displays.
3. Specify where the DHCP snooping bindings database is located.
   - **Local**: The binding table is stored locally on the switch.
   - **Remote**: The binding table is stored on a remote TFTP server. If the database is stored on a remote server, specify the following information:
     - **Remote IP Address**: Specify the IP address of the TFTP server.
     - **Remote File Name**: Specify the file name of the DHCP snooping bindings database in which the bindings are stored.
4. In the **Write Delay** field, specify the time that the switch must wait after writing binding information to persistent storage.
   The delay allows the switch to collect as many entries as possible (new and removed) before writing them to the persistent file. You can specify from 15 to 86400 seconds. By default, the delay is 300 seconds.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Display or clear DHCP snooping statistics

You can display and clear per-interface statistics about the DHCP messages filtered by the DHCP snooping feature on untrusted interfaces.

To display or clear the DHCP snooping statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCP Snooping > Statistics**.
   The Statistics page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Click the **Clear** button to clear all interfaces statistics.
5. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the DHCP snooping statistics.

Table 185.DHCP Snooping Statistics information

Field	Description
MAC Verify Failures	The number of DHCP messages that were dropped because the source MAC address and client hardware address did not match. MAC address verification is performed only if it is globally enabled.
Client Ifc Mismatch	The number of packets that were dropped by DHCP snooping because the interface and VLAN on which the packet was received do not match the client's interface and VLAN information stored in the binding database.
DHCP Server Msgs	The number of DHCP server messages that were dropped on an untrusted port.

DHCPv6 snooping

DHCPv6 snooping is a feature that provides security by filtering untrusted DHCPv6 messages and by building and maintaining a DHCPv6 snooping binding table. An untrusted message is a message that is received from outside the network or firewall and that can cause traffic attacks within your network. The DHCPv6 snooping binding table contains the MAC address, IP address, lease time, binding type, VLAN number, and interface information that corresponds to the local

untrusted interfaces of a switch. An untrusted interface is an interface that is configured to receive messages from outside the network or firewall. A trusted interface is an interface that is configured to receive only messages from within the network.
DHCPv6 snooping acts like a firewall between untrusted hosts and DHCPv6 servers. It also provides a way to differentiate between untrusted interfaces connected to the end user and trusted interfaces connected to the DHCPv6 server or another switch.

Enable DHCPv6 snooping for the switch

You can globally enable DHCPv6 snooping for the switch.

To globally enable DHCPv6 snooping for the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCPv6 Snooping > Global Configuration**.
   The Global Configuration page displays.
3. Select the DHCPv6 Snooping Mode **Enable** radio button.
   The default is Disable.
4. To enable the verification of the sender's MAC address for DHCPv6 snooping, leave the MAC Address Validation **Enable** radio button selected.
   The default is Enable.
   When MAC address validation is enabled, the switch checks packets that are received on an untrusted interface to verify that the MAC address and the DHCPv6 client hardware address match. If the addresses do not match, the switch drops the packet.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Enable DHCPv6 snooping for a VLAN

You can enable DHCPv6 snooping on one or more VLANs.

To enable DHCPv6 snooping for a VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Control > DHCPv6 Snooping > Global Configuration**.
   The Global Configuration page displays.
3. In the **VLAN ID** field, type a VLAN ID.
4. From the **DHCPv6 Snooping Mode** menu, select **Enable**.
5. Click the **Apply** button.
   Your settings are saved.
   The VLAN ID displays in the VLAN Configuration table.
6. To save the settings to the running configuration, click the **Save** icon.

Configure DHCPv6 snooping interface settings

You can display and configure each port as a trusted or untrusted port. Any DHCPv6 responses received on a trusted port are forwarded. If a port is configured as untrusted, any DHCPv6 (or BootP) responses received on that port are discarded.

To configure DHCPv6 snooping interface settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCPv6 Snooping > Interface Configuration**.
   The Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.

4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Trust Mode** menu, select the trust mode:
   - **Disabled**: The interface is considered to be untrusted and could potentially be used to launch a network attack. DHCPv6 server messages are checked against the bindings database. On untrusted ports, DHCPv6 snooping enforces the following security rules:
     - DHCPv6 packets from a DHCPv6 server are dropped.
     - DHCPv6 messages are dropped if the MAC address is in the snooping database but the binding's interface is other than the interface where the message was received.
     - DHCPv6 packets are dropped if the source MAC address does not match the client hardware address and if MAC address validation is globally enabled.
   - **Enabled**: The interface is considered to be trusted and forwards DHCPv6 server messages without validation.
6. From the **Invalid Packets** menu, select the packet logging mode.
   When enabled, the DHCPv6 snooping feature generates a log message when an invalid packet is received and dropped by the interface.
7. In the **Rate Limit (pps)** field, specify the rate limit value for DHCPv6 snooping purposes.
   If the incoming rate of DHCPv6 packets per second exceeds the configured burst interval per second, the port shuts down. If the rate limit is None (which is the default), the burst interval is also not applicable, and rate limiting is disabled.
8. In the **Burst Interval (secs)** field, specify the burst interval in seconds for rate limiting on the interface.
   If the rate limit is N/A, the burst interval is not applicable.
9. Click the **Apply** button.
   Your settings are saved.
10. To save the settings to the running configuration, click the **Save** icon.

Add a static DHCPv6 binding and display or clear dynamic DHCPv6 bindings

You can add a static binding in the DHCPv6 snooping bindings database and display or clear the dynamic bindings in the bindings table.

To add a static DHCPv6 binding and display or clear the dynamic bindings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCPv6 Snooping > Binding Configuration**.
   The Binding Configuration page displays.
3. From the **Interface** menu, select the interface.
4. In the **MAC Address** field, specify the MAC address for the binding to be added.
   This is the key to the binding database.
5. From the **VLAN ID** menu, select the ID of the VLAN.
6. In the **IPv6 Address** field, specify the IPv6 address for the binding to be added.
7. Click the **Add** button.
   The DHCPv6 snooping binding entry is added to the database.
   The Dynamic Binding Configuration table shows information about the DHCPv6 bindings that were learned on each interface on which DHCPv6 snooping is enabled. The following table describes the dynamic binding information.

Field	Description
Interface	The interface on which the DHCPv6 client message was received.
MAC Address	The MAC address associated with the DHCPv6 client that sent the message. This is the key to the binding database.
VLAN ID	The VLAN ID of the client interface.
IPv6 Address	The IPv6 address assigned to the client by the DHCPv6 server.
Lease Time	The remaining IP address lease time for the client.

8. To save the settings to the running configuration, click the **Save** icon.
9. To clear all dynamic bindings from the Dynamic Binding Configuration table, click the **Clear** button.
10. To save the settings to the running configuration, click the **Save** icon.

Remove a static DHCPv6 binding

You can remove a static binding from the DHCPv6 snooping bindings database.

To remove a static DHCPv6 binding:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCPv6 Snooping > Binding Configuration**.
   The Binding Configuration page displays.
3. In the Static Binding Configuration table, select the check box for the static binding.
4. Click the **Delete** button.
   The DHCPv6 snooping binding entry is removed from the database.
5. To save the settings to the running configuration, click the **Save** icon.

Configure DHCPv6 snooping persistent settings

You can configure the persistent location of the DHCPv6 snooping bindings database. The bindings database can be stored locally on the switch or on a remote device in the network. The switch must be able to reach the IP address of the remote device to send bindings to a remote database.

To configure DHCPv6 snooping persistent settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCPv6 Snooping > Persistent Configuration**.
   The Persistent Configuration page displays.
3. Specify where the DHCPv6 snooping bindings database is located.
   - **Local**: The binding table is stored locally on the switch.
   - **Remote**: The binding table is stored on a remote TFTP server. If the database is stored on a remote server, specify the following information:
     - **Remote IP Address**: Specify the IP address of the TFTP server.
     - **Remote File Name**: Specify the file name of the DHCPv6 snooping bindings database in which the bindings are stored.

4. In the **Write Delay** field, specify the time that the switch must wait after writing binding information to persistent storage.
   The delay allows the switch to collect as many entries as possible (new and removed) before writing them to the persistent file. You can specify from 15 to 86400 seconds. By default, the delay is 300 seconds.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Display or clear DHCPv6 snooping statistics

You can display and clear per-interface statistics about the DHCPv6 messages filtered by the DHCPv6 snooping feature on untrusted interfaces.

To display or clear the DHCPv6 snooping statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > DHCPv6 Snooping > Statistics**.
   The Statistics page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Click the **Clear** button to clear all interfaces statistics.
5. To save the settings to the running configuration, click the **Save** icon.

The following table describes the DHCPv6 snooping statistics.

Table 186.DHCPv6 Snooping Statistics information

Field	Description
MAC Verify Failures	The number of DHCPv6 messages that were dropped because the source MAC address and client hardware address did not match. MAC address verification is performed only if it is globally enabled.
Client Ifc Mismatch	The number of packets that were dropped by DHCPv6 snooping because the interface and VLAN on which the packet was received do not match the client's interface and VLAN information stored in the binding database.
DHCPv6 Server Msgs	The number of DHCPv6 server messages that were dropped on an untrusted port.

IP source guard interfaces

You can configure IP source guard (IPSG) on individual interfaces. IPSG is a security feature that filters IP packets based on source ID. This feature helps protect the network from attacks that use IP address spoofing to compromise or overwhelm the network. The source ID can be either the source IP address or a combination of a source IP address and source MAC address, referred to as a pair. The DHCP snooping bindings database, along with IPSG entries in the database, identify authorized source IDs.
If you enable IPSG on a port on which DHCP snooping is disabled or on which DHCP snooping is enabled but the port is untrusted, all IP traffic received on that port is dropped. In addition, IPSG interacts with port security (see Port security on page 618) to enforce the source MAC address in incoming packets. Port security controls how source MAC addresses are learned in the Layer 2 forwarding database (the MAC address table). If a port receives a frame with a previously unlearned source MAC address, port security uses IPSG to determine if the MAC address belongs to a valid binding.

Configure IP source guard on an interface

You can configure IP source guard on individual interfaces.

To configure IP source guard on an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > IP Source Guard > Interface Configuration**.
   The Interface Configuration page displays.

3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **IPSG Mode** menu, specify the administrative mode of IPSG on the interface:
   - **Enable**: IPSG is enabled on the interface. The IP address of the source is not in DHCP snooping binding database, packets are not forwarded.
   - **Disable**: IPSG is disabled on the interface. The default is Disable.
6. From the **IPSG Port Security** menu, specify the administrative mode of IPSG port security on the interface:
   - **Enable**: IPSG port security is enabled on the interface. Packets are not forwarded if the MAC address of the source is not in forwarding database table or the DHCP snooping binding database. To enforce filtering based on the MAC address, you must also configure the following features:
     - Enable port security globally (see Configure the global port security mode on page 619).
     - Enable port security on the interface (see Configure a port security interface on page 620).
   - **Disable**: IPSG port security is disabled on the interface. The default is Disable.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Add a static IP source guard binding and display or clear dynamic IP source guard bindings

You can add a static binding in the IP source guard (IPSG) bindings database and display or clear the dynamic bindings in the bindings table.

To add a static IPSG binding and display or clear the dynamic IPSG bindings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > IP Source Guard > Binding Configuration**.
   The Binding Configuration page displays.
3. From the **Interface** menu, select the interface.
4. In the **MAC Address** field, specify the MAC address for the binding to be added.
   This is the key to the binding database.
5. From the **VLAN ID** menu, select the ID of the VLAN.
6. In the **IP Address** field, specify the IP address for the binding to be added.
7. Click the **Add** button.
   The IPSG static binding entry is added to the database.
   The Filter Type field displays if IPSG is configured to function with the source IP address only or, if port security is also enabled for IPSG, with both the source IP address and MAC address.
   The Dynamic Binding Configuration table shows information about the IPSG bindings that were learned on each interface on which IPSG is enabled. The following table describes the dynamic binding information.

Field	Description
Interface	The interface on which the IPSG source binding was learned.
MAC Address	The MAC address associated with the IPSG source device. This is the key to the binding database.
VLAN ID	The VLAN ID for the interface of the IPSG source device.
IP Address	The IP address of the IPSG source device.
Filter Type	The filter type used on the switch interface. Either IPSG is configured to function with the source IP address only or, if port security is also enabled for IPSG, with both the source IP address and MAC address.

8. To save the settings to the running configuration, click the **Save** icon.

9. To clear all dynamic bindings from the Dynamic Binding Configuration table, click the **Clear** button.
10. To save the settings to the running configuration, click the **Save** icon.

Remove a static IP source guard binding

You can remove a static binding from the IP source guard (IPSG) bindings database.

To remove a static IPSG binding:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > IP Source Guard > Binding Configuration**.
   The Binding Configuration page displays.
3. In the Static Binding Configuration table, select the check box for the static binding.
4. Click the **Delete** button.
   The IPSG binding entry is removed from the database.
5. To save the settings to the running configuration, click the **Save** icon.

IPv6 source guard interfaces

You can configure IPv6 source guard (IPv6SG) on individual interfaces. IPv6SG is a security feature that filters IPv6 packets based on source ID. This feature helps protect the network from attacks that use IPv6 address spoofing to compromise or overwhelm the network. The source ID can be either the source IPv6 address or a combination of a source IPv6 address and source MAC address, referred to as a pair. The DHCP snooping bindings database, along with IPv6SG entries in the database, identify authorized source IDs.
If you enable IPv6SG on a port on which DHCP snooping is disabled or on which DHCP snooping is enabled but the port is untrusted, all IPv6 traffic received on that port is dropped. In addition, IPv6SG interacts with port security (see Port security on page 618) to enforce the source MAC address in incoming packets. Port security controls how source MAC addresses are learned in the Layer 2 forwarding database (the MAC address table). If a port receives a frame with a previously unlearned source MAC address, port security uses IPv6SG to determine if the MAC address belongs to a valid binding.

Configure IPv6 source guard on an interface

You can configure IPv6 source guard on individual interfaces.

To configure IPv6 source guard on an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > IPv6 Source Guard > Interface Configuration**.
   The Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **IPv6SG Mode** menu, specify the administrative mode of IPv6SG on the interface:
   - **Enable**: IPv6SG is enabled on the interface. The IPv6 address of the source is not in DHCP snooping binding database, packets are not forwarded.
   - **Disable**: IPv6SG is disabled on the interface. The default is Disable.

6. From the **IPv6SG Port Security** menu, specify the administrative mode of IPv6SG port security on the interface:
   - **Enable**: IPv6SG port security is enabled on the interface. Packets are not forwarded if the MAC address of the source is not in forwarding database table or the DHCP snooping binding database. To enforce filtering based on the MAC address, you must also configure the following features:
     - Enable port security globally (see Configure the global port security mode on page 619).
     - Enable port security on the interface (see Configure a port security interface on page 620).
   - **Disable**: IPv6SG port security is disabled on the interface. The default is Disable.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

Add a static IPv6 source guard binding and display or clear dynamic IPv6 source guard bindings

You can add a static binding in the IPv6 source guard (IPv6SG) bindings database and display or clear the dynamic bindings in the bindings table.

To add a static IPv6SG binding and display or clear the dynamic IPv6SG bindings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > IPv6 Source Guard > Binding Configuration**.
   The Binding Configuration page displays.
3. From the **Interface** menu, select the interface.
4. In the **MAC Address** field, specify the MAC address for the binding to be added.
   This is the key to the binding database.
5. From the **VLAN ID** menu, select the ID of the VLAN.
6. In the **IPv6 Address** field, specify the IPv6 address for the binding to be added.

7. Click the **Add** button.
   The IPv6SG static binding entry is added to the database.
   The Filter Type field displays if IPv6SG is configured to function with the source IPv6 address only or, if port security is also enabled for IPv6SG, with both the source IPv6 address and MAC address.
   The Dynamic Binding Configuration table shows information about the IPv6SG bindings that were learned on each interface on which IPv6SG is enabled. The following table describes the dynamic binding information.

Field	Description
Interface	The interface on which the IPv6SG source binding was learned.
MAC Address	The MAC address associated with the IPv6SG source device. This is the key to the binding database.
VLAN ID	The VLAN ID for the interface of the IPv6SG source device.
IPv6 Address	The IPv6 address of the IPv6SG source device.
Filter Type	The filter type used on the switch interface. Either IPv6SG is configured to function with the source IPv6 address only or, if port security is also enabled for IPv6SG, with both the source IPv6 address and MAC address.

8. To save the settings to the running configuration, click the **Save** icon.
9. To clear all dynamic bindings from the Dynamic Binding Configuration table, click the **Clear** button.
10. To save the settings to the running configuration, click the **Save** icon.

Remove a static IPv6 source guard binding

You can remove a static binding from the IPv6 source guard (IPv6SG) bindings database.

To remove a static IPv6SG binding:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > IPv6 Source Guard > Binding Configuration**.
   The Binding Configuration page displays.
3. In the Static Binding Configuration table, select the check box for the static binding.
4. Click the **Delete** button.
   The IPv6SG binding entry is removed from the database.

5. To save the settings to the running configuration, click the **Save** icon.

Dynamic ARP inspection

Dynamic ARP inspection (DAI) is a security feature that rejects invalid and malicious ARP packets. The feature prevents a class of man-in-the-middle attacks, where an unfriendly station intercepts traffic for other stations by poisoning the ARP caches of its unsuspecting neighbors. The unfriendly station sends ARP requests or responses mapping another station's IP address to its own MAC address.

DAI relies on DHCP snooping. DHCP snooping listens to DHCP message exchanges and builds a bindings database of valid MAC addresses, IP addresses, VLAN interfaces, and so on.

If DAI is enabled and if a sender MAC address and sender IP address do not match an entry in the DHCP snooping bindings database, the switch drops the ARP packet if. However, you can also create static mappings in the DHCP snooping bindings database. Static mappings are useful when hosts configure static IP addresses, the switch cannot run DHCP snooping, or other switches in the network do not run dynamic ARP inspection. A static mapping associates an IP address to a MAC address on a VLAN.

You can configure DAI VLANs, interfaces, and access control lists (ACLs) with associated rules.

Configure the global DAI settings

You can configure the global dynamic ARP inspection (DAI) settings.

To configure the global DAI settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Dynamic ARP Inspection > DAI Configuration**.
   The DAI Configuration page displays.
3. Select the Validate Source MAC **Disable** or **Enable** radio button.
   This settings specifies the DAI source MAC validation mode for the switch. If you select Enable, the sender's MAC address in an ARP packet is validated. The default is Disable.
4. Select the Validate Destination MAC **Disable** or **Enable** radio button
   This setting specifies the DAI destination MAC validation mode for the switch. If you select Enable, the destination MAC address in an ARP response packet is validated. The default is Disable.
5. Select the Validate IP **Disable** or **Enable** radio button.
   This setting specifies the DAI IP validation mode for the switch. If you select Enable, the IP address in an ARP packet is validated. The default is Disable.

6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Configure DAI VLANs

You can configure one or more dynamic ARP inspection (DAI) VLANs.

To configure one or more DAI VLANs:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Dynamic ARP Inspection > DAI VLAN Configuration**.
   The DAI VLAN Configuration page displays.
   The table lists the VLANs.
3. Select the check box for the VLAN.
4. From the **Admin Mode** menu, select **Enable** or **Disable** to specify if DAI is enabled for the VLAN.
   The default is Disable.
5. From the **Invalid Packets** menu, select **Enable** or **Disable** to specify if DAI logging is enabled for the VLAN.
   The default is Enable.
6. In the **ARP ACL Name** field, specify a name of an existing ARP ACL (see Create a DAI access control list on page 669).
   The ARP ACL is used for ARP packet validation. To remove an existing ARP ACL name from the **ARP ACL Name** field, enter **N/A**.
7. From the **Static Flag** menu, select **Enable** or **Disable** to specify if the ARP packet must be validated using the DHCP snooping database if the ARP ACL rule does not match.
   If enabled, the ARP packet is validated by the ARP ACL rule only. If disabled, the ARP packet needs further validation by using the DHCP snooping database. The default is Disable.
8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Configure DAI interfaces

You can configure one or more dynamic ARP inspection (DAI) interfaces.

To configure one or more DAI interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Dynamic ARP Inspection > DAI Interface Configuration**.
   The DAI Interface Configuration page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.
4. Select one or more interfaces by taking one of the following actions:
   - To configure a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
   - To configure multiple interfaces with the same settings, select the check box associated with each interface.
   - To configure all interfaces with the same settings, select the check box in the heading row.
5. From the **Trust Mode** menu, specify if the interface is trusted for DAI.
   If enabled, the interface is trusted and ARP packets entering the interface are forwarded without checking. If disabled, the interface is not trusted and ARP packets entering the interface are subjected to DAI. The default is Disable.
6. In the **Rate Limit (pps)** field, specify the rate limit value for DAI.
   If the incoming rate of ARP packets exceeds the specified value for consecutive burst interval seconds, ARP packets are dropped. If you specify N/A, no limit exists. The range is from 0 to 300. The default is 15 packets per second (pps).
7. In the **Burst Interval (secs)** field, specify the burst interval value for rate limiting on the interface. If you specify N/A, the burst interval is not effective. The range is from 1 to 15. The default is 1 second.

8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

Create a DAI access control list

You can create a dynamic ARP inspection (DAI) access control list (ACL) to which you then can add rules.

To create a DAI ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Dynamic ARP Inspection > DAI ACL Configuration**.
   The DAI ACL Configuration page displays.
3. In the **Name** field, specify a name of up to 31 characters.
4. Click the **Add** button.
   The DAI ACL is added.
5. To save the settings to the running configuration, click the **Save** icon.

Configure a rule for an existing DAI ACL

After you add a DAI ACL (see Create a DAI access control list on page 669), you can configure a rule for it.

To configure a rule for an existing DAI ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Dynamic ARP Inspection > DAI ACL Rule Configuration**.
   The DAI ACL Rule Configuration page displays.
3. From the **ACL Name** menu, select the DAI ACL for which you want to configure the rule.
4. From the **Action** menu, select the action that must occur if a packet matches the rule's criteria:
   - **Permit**: The packet is permitted and forwarded.
   - **Deny**: The packet is denied and dropped.

5. In the **Source IP Address** field, enter the source IP address that must be used as a match for the rule.
6. In the **Source MAC Address** field, enter the source MAC address that must be used as a match for the rule.
7. Click the **Add** button.
   The rule is added.
8. To save the settings to the running configuration, click the **Save** icon.

Delete a rule from an existing DAI ACL

You can delete a rule from an existing DAI ACL.

To delete a rule from an existing DAI ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Dynamic ARP Inspection > DAI ACL Rule Configuration**.
   The DAI ACL Rule Configuration page displays.
3. From the **ACL Name** menu, select the DAI ACL from which you want to delete a rule.
4. Select the check box for the rule.
5. Click the **Delete** button.
   The rule is removed.
6. To save the settings to the running configuration, click the **Save** icon.

Delete a DAI access control list

You can delete a dynamic ARP inspection (DAI) access control list (ACL) that you no longer need.

To delete a DAI ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Dynamic ARP Inspection > DAI ACL Configuration**.
   The DAI ACL Configuration page displays.
3. Select the check box for the ACL.

4. Click the **Delete** button.
   The DAI ACL is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Display the DAI statistics

To display the DAI statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Dynamic ARP Inspection > DAI Statistics**.
   The DAI Statistic page displays.
3. To refresh the page, click the **Refresh** button.
4. To clear the DAI statistics, click the **Clear** button.
5. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 187.DAI Statistics information

Field	Description
VLAN	The VLAN ID.
DHCP Drops	The number of ARP packets that were dropped by DAI because no matching DHCP snooping binding entry exists.
DHCP Permits	The number of ARP packets that were forwarded by DAI because a matching DHCP snooping binding entry exists.
ACL Drops	The number of ARP packets that were dropped by DAI because no matching ARP ACL rule exists for the VLAN and the static flag is set on the VLAN.
ACL Permits	The number of ARP packets that were permitted by DAI because a matching ARP ACL rule exists for the VLAN.
Bad Source MAC	The number of ARP packets that were dropped by DAI because the sender MAC address in the ARP packets did not match the source MAC address in the Ethernet header.
Bad Dest MAC	The number of ARP packets that were dropped by DAI because the target MAC address in the ARP reply packets did not match the destination MAC address in the Ethernet header.
Invalid IP	The number of ARP packets that were dropped by DAI because the sender IP address in the ARP packets or the target IP address in the ARP reply packets is invalid. Invalid addresses include 0.0.0.0, 255.255.255.255, IP multicast addresses, class E addresses (240.0.0.0/4), and loopback addresses (127.0.0.0/8).
Forwarded	The number of valid ARP packets forwarded by DAI.
Dropped	The number of invalid ARP packets dropped by DAI.

Captive portals

The captive portal feature allows you to prevent clients from accessing the network until user verification is established. You can configure captive portal verification to allow access for both guest and authenticated users. Authenticated users must be validated against a database of authorized captive portal users before access is granted. The database can be stored locally on the switch or on a RADIUS server.

The authentication server supports both HTTP and HTTPS web connections. In addition, you can configure a captive portal to use an optional HTTP port (in support of HTTP proxy networks). If configured, this additional port is then used exclusively by the captive portal. This optional port is

in addition to the standard HTTP port 80, which is used for all other web traffic.
If you enable the captive portal feature on a port, the port drops all traffic from unauthenticated clients except for ARP, DHCP, DNS, and NETBIOS packets, which are forwarded so that unauthenticated clients can get an IP address and resolve the host name or domain names. Data traffic from authenticated clients goes through, and the rules do not apply to these packets.
For a port on which you enable the captive portal feature, if an unauthenticated client opens a web browser and tries to connect to network, the captive portal redirects all HTTP and HTTPS traffic from unauthenticated clients to the authenticating server on the switch. A captive portal web page is displayed for the unauthenticated client, allowing the client to authenticate, after which the client receives access to the port.
The captive portal feature is not supported for VLAN interfaces, loopback interfaces, and logical interfaces. The captive portal feature uses MAC-address based authentication and not port-based authentication. This means that all clients connected to the captive portal interface must be authenticated before they can get access to the network.

Configure the global captive portal settings

You can control the administrative state of the captive portal feature, and configure global settings that affect all captive portals that you configure on the switch.

To configure the global captive portal settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Captive Portal > CP Global Configuration**.
   The CP Global Configuration page displays.
3. Select the Admin Mode **Disable** or **Enable** radio button.
   This selection sets the administrative mode of the captive portal feature. The default is Disable.
4. In the **Additional HTTP Port** field, enter a port number between 0 and 65535, excluding port 80.
   HTTP traffic uses standard port 80, but you can configure an additional port for HTTP traffic. The default is 0, which disables this option.
5. In the **Additional HTTP Secure Port** field, enter a port number between 0 and 65535, excluding port 443.
   HTTP Secure (HTTPS) traffic uses standard port 443, but you can configure an additional port for HTTPS traffic. The default is 0, which disables this option.

6. In the **Authentication Timeout** field, enter the number of seconds that an authentication session remains open while a client attempts to access the network through a captive portal. To access the network through a captive portal, the client must first enter authentication information on an authentication web page. When the authentication session time-out expires and the client is not yet authenticated, the switch disconnects any active TCP or SSL connection with the client. The range is from 60 to 600 seconds. The default is 300 seconds.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 188.Captive portal global configuration information

Field	Description
Operational Status	The operational status of the captive portal feature, which is either Enabled or Disabled.
Disabled Reason	If the captive portal feature is disabled, the reason that it is disabled: • Administrator Disabled • IP Address Not Configured • No IP Routing Interface • Routing Disabled
CP IP Address	The IP address that the captive portal uses.
Supported Captive Portals	The total number of captive portals that are supported on the switch.
Configured Captive Portals	The number of captive portals that are configured on the switch.
Active Captive Portals	The number of captive portal that are operationally enabled (active).
System Supported Users	The total number of users that can be authenticated locally and remotely on the switch.
Local Supported Users	The total number of users that can be authenticated locally on the switch.
Configured Local Users	The number of local users that are configured on the switch.
Authenticated Users	The number of users that are authenticated for all captive portal on the switch.

Configure a captive portal

By default, the switch provides one default captive portal. You can change the settings for the default captive portal. You can also add and configure up to nine additional portals.

To configure a captive portal instance:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Captive Portal > CP Configuration**.
   The CP Configuration page displays.
3. To do one of the following:
   - **Change the default captive portal**: Next to the **CP ID** field, select the **1** check box.
   - **Add a new captive portal**: From the **CP ID** menu, select an ID from **2** to **10**.
   - **Change a captive portal that you added before**: Next to the **CP ID** field, select the check box that is associated with the captive portal.
4. In the **CP Name** field, specify a name for the captive portal.
   The name can be up to 31 alphanumeric characters.
5. From the **Admin Mode** menu, select **Enable** or **Disable**.
   This selection sets the administrative mode of the captive portal. By default, a captive portal is disabled.
6. In the **Protocol** field, select **HTTP** or **HTTPS** used for communication with clients during the verification process:
   - **HTTP**: Does not use encryption during verification.
   - **HTTPS**: Uses the Secure Sockets Layer (SSL), which requires a certificate for encryption. The certificate is presented to the user at connection time.
7. From the **Verification** menu, select the type of user authentication that the captive portal uses when a client attempt to connect:
   - **Guest**: No user authentication is required.
   - **Local**: The user is authenticated through the local database.
   - **RADIUS**: The user is authenticated through a RADIUS server.
8. Select the block status, which is a way to quickly and temporarily control a captive portal during unexpected events, such as denial of service attacks:
   - **Disable**: The blocking function is disabled. You cannot quickly block the captive portal.
   - **Enable**: The captive portal is blocked.
   - **Not Blocked**: The captive portal is not blocked. The default is Not Blocked.

9. From the **Group** menu, select an existing captive portal user group for the captive portal. All users who belong to the group are permitted to access the network through the captive portal.
   This option is available only if you select **Local** or **RADIUS** from the **Verification** menu.
10. In the **Idle Timeout** field, enter the number of seconds that the switch waits before terminating an idle session and logging out the user.
    The range is from 0 to 900 seconds. The default is 0 seconds, which means that the time-out is not enforced.
11. From the **User Logout** menu, specify if an authenticated client can log out from the captive portal and network:
    - **Enable**: A user can deliberately log out from the captive portal and network.
    - **Disable**: A user cannot log out from the captive portal and network, causing the connection status to remain authenticated until the captive portal deauthenticates the user, for example because idle time-out or session time-out period is exceeded.
12. If the verification mode is RADIUS (that is, you select **RADIUS** from the **Verification** menu), specify the IP address of the RADIUS server in the **Radius Auth Server** field.
13. From the **Redirect Mode** menu, specify if newly authenticated clients are redirected to a URL:
    - **Enable**: Authenticated clients are redirected to a URL that you can specify in the Redirect URL field.
    - **Disable**: Authenticated clients are not redirected and the default locale “welcome” is used. The default is Disable.
14. In the **Redirect URL** field, specify the URL to which newly authenticated clients must be redirected.
    The maximum length for the URL is 512 alphanumeric characters. The default URL is http://*<ipaddress>*/v1/security/captive_portal/cp_welcome.html.
    This option is available only if you select **Enable** from the **Redirect Mode** menu.
15. In the **Background Color** field, specify the value of the background color.
    The default color is #BFBFBF.
16. In the **Foreground Color** field, specify the value of the foreground color.
    The default color is #999999.
17. In the **Separator Color** field, specify the value of the separator color.
    The default color is #B70024.

18. In the **Max Bandwidth Down** field, specify the maximum rate at which a client can receive data from the network.
    The rate is in bytes per seconds. The range is from 0 to 536870911. The default is 0, which indicates that the limit is not enforced.
19. In the **Max Bandwidth Up** field, specify the maximum rate at which a client can send data into the network.
    The rate is in bytes per seconds. The range is from 0 to 536870911. The default is 0, which indicates that the limit is not enforced.
20. In the **Max Input** field, specify the maximum number of octets that a client is allowed to receive.
    After this limit is reached, the client is disconnected. The range is from 0 to 4294967295. The default is 0, which indicates that the limit is not enforced.
21. In the **Max Output** field, specify the maximum number of octets that a client is allowed to transmit.
    After this limit is reached, the client is disconnected. The range is from 0 to 4294967295. The default is 0, which indicates that the limit is not enforced.
22. In the **Max Total** field, specify the maximum number of octets that a client is allowed to transfer, meaning the sum of octets transmitted and received.
    After this limit is reached, the client is disconnected. The range is from 0 to 4294967295. The default is 0, which indicates that the limit is not enforced.
23. Do one of the following:
    - If you are adding a new captive portal, click the **Add** button. Your settings are saved. The captive portal is added.
    - If you are changing the default captive portal or a captive portal that you added before, click the **Apply** button. Your settings are saved.
24. To save the settings to the running configuration, click the **Save** icon.

Delete a captive portal

You can delete a captive portal that you no longer need.

To delete a captive portal:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Captive Portal > CP Configuration**.
   The CP Configuration page displays.

3. Select the check box for the captive portal.
4. Click the **Delete** button.
   The captive portal is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure a captive portal binding

You can associate a captive portal with one or more interfaces. Although you can associate multiple interfaces with one captive portal, you can associate an interface with one captive portal only.

To configure a captive portal binding:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Captive Portal > CP Binding Configuration**.
   The CP Binding Configuration page displays.
3. From the **CP ID** menu, select the captive portal.
   The CP Name field displays the captive portal name.
4. To add the selected captive portal to a port, in the Ports table (or if a stack is configured, in one of the Ports tables), click the port so that a check mark displays.
   You can add the captive portal to several ports.
   The ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed, and you can select ports on different stacked switches.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Display or delete captive portal bindings in the captive portal binding table

You can display or delete bindings for captive portals.

To display or delete captive portal bindings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Control > Captive Portal > CP Binding Table**.
   The CP Binding Table page displays. This table displays the captive portal-to-interface bindings.
3. To delete a captive portal-to-interface binding, do the following:
   a. Select the check box next to the interface.
   b. Click the **Delete** button.
      The binding is removed.
4. To save the settings to the running configuration, click the **Save** icon.
5. To refresh the page, click the **Refresh** button.
   The following table describes the view-only fields on the page.

Table 189.Captive portal binding information

Field	Description
Interface	The interface to which the captive portal is bound.
CP ID	The ID of the captive portal.
Operational Status	Indicates if the portal is active on the interface (Disable or Enable).
Block Status	Indicates if the captive portal is temporarily blocked for authentication.
Authenticated Users	The number of users authenticated through the captive portal on the interface.

Configure captive portal groups

Captive portal groups let you organize captive portal users, which is useful if a large number of users must be able to access a captive portal. After you add a captive portal group, you can add users to the group.
By default, all captive portal users are added to the default group with ID 1. You can add up to 10 groups.

Add a captive portal group

After you add a captive portal group, you can add captive portal users to the group (see Configure captive portal users on page 680).

To add a captive portal group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > Control > Captive Portal > CP Group Configuration**.
   The CP Group Configuration page displays.
3. From the **Group ID** menu, select an ID for the group.
   You can select an ID from **1** to **10**.
4. In the **Group Name** field, specify a name for the group.
   The name can contain from 1 to 31 alphanumeric characters.
5. Click the **Add** button.
   The group is added.
6. To save the settings to the running configuration, click the **Save** icon.

Remove a captive portal group

You can remove a captive portal group that you no longer need.

To remove a captive portal group:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Captive Portal > CP Group Configuration**.
   The CP Group Configuration page displays.
3. From the **Group ID** menu, select an ID for the group.
   You can select an ID from **1** to **10**.
4. Select the check box for the group that you want to remove.
5. Click the **Delete** button.
   The group is removed.
6. To save the settings to the running configuration, click the **Save** icon.

Configure captive portal users

You can configure the settings for captive portal users and add them to a captive portal group. By default, all captive users are added to the default group with ID 1. After you add a custom captive portal group (see Configure captive portal groups on page 679), you can also assign captive portal users to the custom group.
The settings that you can configure for an individual captive portal user include a password, time-out periods, bandwidths, and maximum traffic amounts.

Add or modify a captive portal user account

You can add a captive portal user account or modify an existing one.

To add or modify a captive portal user account:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Captive Portal > CP User Configuration**.
   The CP User Configuration page displays.
3. If you change the settings for an existing user account, select the check box for the user account.
   If you change the settings, you cannot change the user ID and user name.
4. If you are adding a user account, in the **User ID** field, specify an ID from **1** to **99** to identify the user.
   After you add the user account, you cannot change the ID.
5. If you are adding a user account, in the **User Name** field, enter a name for the user account. The name can contain up to 31 alphanumeric characters. After you add the user account, you cannot change the name.
6. From the **Edit Password** menu, select to use a password or no password at all for the new user account:
   - **Enable**: The Password and Confirm Password fields are unmasked so that you can specify a unique password for the new user account:
     - In the **Password** field, specify a password from 8 to 64 characters.
     - In the **Confirm Password** field, specify the password again.
   - **Disable**: The Password and Confirm Password fields are masked and you cannot specify a password for the new user account.

   By default, the selection is Disable, which means that no password is used and the password fields are masked out.
7. From the **Group** menu, select at least one captive portal group.
   To assign a user account to more than one group, press the **Control** key on your keyboard and select each group. By default, a new user account is assigned to the default group with ID 1.

8. In the **Session Timeout** field, enter the period in seconds that a user is permitted to remain connected to the network.
   After the session time-out period is exceeded, the user is logged out automatically. The default is 0, which means that the session does not time out. The range is from 0 to 86400 seconds.
9. In the **Idle Timeout** field, enter the period in seconds that the switch waits before terminating an idle session and logging out the user.
   The default is 0, which means that an idle session does not time out. The range is from 0 to 900 seconds.
10. In the **Max Bandwidth Down** field, enter the maximum rate, in bits per second, at which a client can receive data from the network.
    A value of 0 indicates that the global configuration must be used. The range is from 0 to 536870911 bps.
11. In the **Max Bandwidth Up** field, enter the maximum rate, in bits per second, at which a client can send data into the network.
    A value of 0 indicates that the global configuration must be used. The range is from 0 to 536870911 bps.
12. In the **Max Input** field, enter the number of octets the user is allowed to receive.
    After this limit is reached, the user is disconnected. A value of 0 indicates that the global limit must be used. The range is from 0 to 4294967295.
13. In the **Max Output** field, enter the number of octets the user is allowed to transmit.
    After this limit is reached, the user is disconnected. A value of 0 indicates that the global limit must be used. The range is from 0 to 4294967295.
14. In the **Max Total** field, enter the number of bytes the user is allowed to transmit and receive.
    The maximum number of octets is the sum of octets transmitted and received. After this limit is reached, the user is disconnected. A value of 0 indicates that the global limit must be used. The range is from 0 to 4294967295.
15. Do one of the following:
    - If you are adding a new user account, click the **Add** button. Your settings are saved. The user account is added.
    - If you are changing the settings for a user account, click the **Apply** button. Your settings are saved.
16. To save the settings to the running configuration, click the **Save** icon.

Delete a captive portal user account

You can delete a captive portal user account that is no longer required.

To delete a captive portal user account:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Captive Portal > CP User Configuration**.
   The CP User Configuration page displays.
3. Select the check box for the user account.
4. Click the **Delete** button.
   The user account is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure the captive portal trap flag settings

You can specify if the switch sends SNMP traps for actions that occur on the captive portal. By default, the switch does not send any traps for actions that occur on the captive portal.

To configure the captive portal trap flag settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Captive Portal > CP Trap Flags**.
   The CP Trap Flags page displays.
3. Configure the following trap flag settings:
   - **CP Trap Mode**: Enable or disable the captive portal trap mode.
   - **Client Auth Failure**: Enable or disable the transmission of a trap if a client fails to authenticate with a captive portal.
   - **Client Connect**: Enable or disable the transmission of a trap if a client authenticates with and connects to a captive portal.
   - **Client DB Full**: Enable or disable the transmission of a trap each time an entry cannot be added to the client database because it is full.
   - **Client Disconnect**: Enable or disable the transmission of a trap if a client disconnects from a captive portal.

   By default, these trap flag settings are disabled.

4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Display or clear captive portal client statistics

You can display or clear captive portal client statistics, including statistics about traffic sent and received by individual clients.

To display or clear the captive portal client statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Control > Captive Portal > CP Client**.
   The CP Client page displays.
3. To search for a captive portal user by MAC address, specify the MAC address, and click the **Go** button.
   If the entry exists, the information for the captive portal user is displayed. An exact match is required.
4. To refresh the page, click the **Refresh** button.
5. To clear the captive portal client statistics, click the **Clear** button.
6. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 190.Captive portal client statistics

Field	Description
MAC Address	The client MAC address.
IP Address Drops	The IP address of the client.
Protocol	The connection protocol, which is either HTTP or HTTPS.
Verification	The authentication type, which is Guest, Local, or RADIUS.
Session Time	The period that passed since the client was authorized.
Interface	The interface the client is using.
CP ID	The ID of the captive portal.
User Name	The user name (or guest ID) of the connected client.
Bytes Received	The total number of bytes the client received.
Bytes Transmitted	The total number of bytes the client transmitted.
Packets Received	The total number of packets the client received.
Packets Transmitted	The total number of packets the client transmitted.

Access control lists

Access control lists (ACLs) ensure that only authorized users can access specific resources while blocking any unwarranted attempts to reach network resources. ACLs are used to provide traffic flow control, restrict contents, decide which types of traffic are forwarded or blocked, and provide security for the network. The switch supports a total of 100 ACLs, which can be a combination of MAC ACLs, basic IPv4 ACL, extended IPv4 ACLs, and IPv6 ACLs.

To configure an ACL:

1. Create an IPv4-based, IPv6-based, or MAC-based ACL ID.
2. Create a rule and assign it to a unique ACL ID.
3. Define the rules, which can identify protocols, source, and destination IP and MAC addresses, and other packet-matching criteria.
4. Use the ID number to assign the ACL to a port or to a LAG.
   To view ACL configuration examples, see Access control lists (ACLs) on page 793.

Use the ACL Wizard to create a simple ACL

The ACL Wizard helps you create a simple ACL and apply it to the selected ports easily and quickly. First, select an ACL type to use when you create an ACL. Then add an ACL rule to this ACL and apply this ACL on the selected ports.

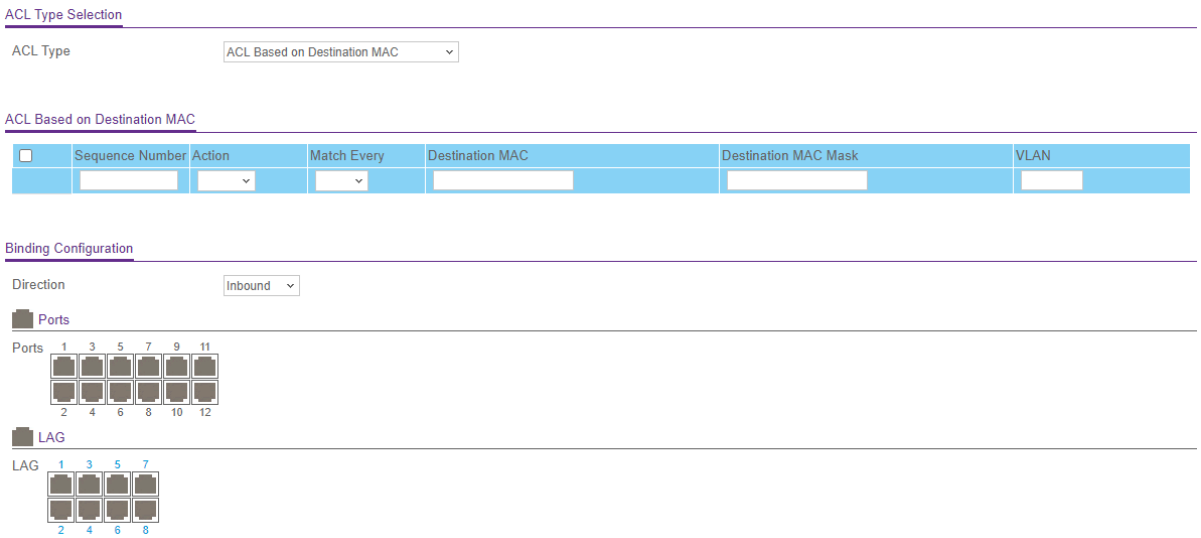
NOTE: The steps in the following procedure describe how you can create an ACL based on the destination MAC address. If you select a different type of ACL (or example, an ACL based on a source IPv4), the page displays different information.

Use the ACL Wizard to create an ACL

The ACL Wizard provides a way to automatically create an ACL based on a type of ACL that you must select and information that you must provide.

To use the ACL Wizard to create an ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > ACL Wizard**.



NOTE: The steps in this procedure describe creating an ACL based on the destination MAC address. If you select a different ACL type, for example, ACL based on a source IPv4, then what is shown on this page varies, depending on the current step in the rule configuration process.

3. From the **ACL Type** menu, select the type of ACL.
   You can select from the following ACL types:
   - **ACL Based on Destination MAC**: Creates an ACL based on the destination MAC address, destination MAC mask, and VLAN.
   - **ACL Based on Source MAC**: Creates an ACL based on the source MAC address, source MAC mask, and VLAN.
   - **ACL Based on Destination IPv4**: Creates an ACL based on the destination IPv4 address and IPv4 address mask.
   - **ACL Based on Source IPv4**: Creates an ACL based on the source IPv4 address and IPv4 address mask.
   - **ACL Based on Destination IPv6**: Creates an ACL based on the destination IPv6 prefix and IPv6 prefix length.
   - **ACL Based on Source IPv6**: Creates an ACL based on the source IPv6 prefix and IPv6 prefix length.
   - **ACL Based on Destination IPv4 L4 Port**: Creates an ACL based on the destination IPv4 Layer 4 port number.
   - **ACL Based on Source IPv4 L4 Port**: Creates an ACL based on the source IPv4 Layer 4 port number.
   - **ACL Based on Destination IPv6 L4 Port**: Creates an ACL based on the destination IPv6 Layer 4 port number.
   - **ACL Based on Source IPv6 L4 Port**: Creates an ACL based on the source IPv6 Layer 4 port number.

   NOTE: For L4 port options, two rules are created (one for TCP and one for UDP).

4. In the **Sequence Number** field, enter a whole number in the range from 1 to 2147483647 that is used to identify the rule.

5. From the **Action** menu, select **Permit** or **Deny** to specify the action that must be taken if a packet matches the rule's criteria.

6. From the **Match Every** menu, select one of the following options:
   - **False**: Packets do not need to match the selected ACL and rule. With this selection, you can add a destination MAC address, destination MAC mask, and VLAN.
   - **True**: All packets must match the selected ACL and rule and are either permitted or denied. In this case, since all packets match the rule, the option of configuring other match criteria is not offered.

7. Specify the additional match criteria for the selected ACL type.
   The rest of the rule match criteria fields available for configuration depend on the selected ACL type. For information about the possible match criteria fields, see the following table.

ACL Based On	Fields
Destination MAC	○ **Destination MAC**: Specify the destination MAC address to compare against an Ethernet frame. The format is xx:xx:xx:xx:xx:xx. The BPDU keyword might be specified using a destination MAC address of 01:80:C2:xx:xx:xx. ○ **Destination MAC Mask**: Specify the destination MAC address mask, which represents the bits in the destination MAC address to compare against an Ethernet frame. The format is xx:xx:xx:xx:xx:xx. The BPDU keyword might be specified using a destination MAC mask of 00:00:00:ff:ff:ff. ○ **VLAN**: Specify the VLAN ID to match within the Ethernet frame.
Source MAC	○ **Source MAC**: Specify the source MAC address to compare against an Ethernet frame. The format is xx:xx:xx:xx:xx:xx. ○ **Source MAC Mask**: Specify the source MAC address mask, which represents the bits in the source MAC address to compare against an Ethernet frame. The format is (xx:xx:xx:xx:xx:xx). ○ **VLAN**: Specify the VLAN ID to match within the Ethernet frame.
Destination IPv4	○ **Destination IP Address**: Specify the destination IP address. ○ **Destination IP Mask**: Specify the destination IP address mask.
Source IPv4	○ **Source IP Address**: Specify the source IP address. ○ **Source IP Mask**: Specify the source IP address mask.
Destination IPv6	○ **Destination Prefix**: Specify the destination prefix. ○ **Destination Prefix Length**: Specify the destination prefix length.
Source IPv6	○ **Source Prefix**: Specify the source destination prefix. ○ **Source Prefix Length**: Specify the source prefix length.
Destination IPv4 L4 Port	○ **Destination L4 port (protocol)**: Specify the destination IPv4 L4 port protocol. ○ **Destination L4 port (value)**: Specify the destination IPv4 L4 port value.
Source IPv4 L4 Port	○ **Source L4 port (protocol)**: Specify the source IPv4 L4 port protocol. ○ **Source L4 port (value)**: Specify the source IPv4 L4 port value.
Destination IPv6 L4 Port	○ **Destination L4 port (protocol)**: Specify the destination IPv6 L4 port protocol. ○ **Destination L4 port (value)**: Specify the destination IPv6 L4 port

ACL Based On	Fields
	value.
Source IPv6 L4 Port	○ **Source L4 port (protocol)**: Specify the source IPv6 L4 port protocol. ○ **Source L4 port (value)**: Specify the source IPv6 L4 port value.

8. As a sample, the following steps describe how you can create an ACL based on the destination MAC address:
   a. In the **Destination MAC** field, specify the destination MAC address that must be compared against the information in an Ethernet frame.
   The format is xx:xx:xx:xx:xx:xx. The BPDU keyword can be specified using a destination MAC address of 01:80:C2:xx:xx:xx.
   b. In the **Destination MAC Mask** field, specify the destination MAC address mask that must be compared against the information in an Ethernet frame.
   The format is xx:xx:xx:xx:xx:xx. The BPDU keyword can be specified using a destination MAC mask of 00:00:00:ff:ff:ff.
   c. In the **VLAN ID** field, specify which VLAN must be compared against the information in an Ethernet frame.
   The range is from 1 to 4093. Either a VLAN range or VLAN can be configured.
9. In the Binding Configuration section, from the **Direction** menu, select the packet filtering direction for the ACL.
   - **Inbound**: The MAC binding can be applied to incoming traffic only.
   - **Outbound**: The MAC binding can be applied to outgoing traffic only.
10. In the Ports and LAG tables in the Binding Configuration section, select the ports and LAGs to which the ACL must be applied.
    The ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed. In addition, all LAGs are displayed, whether or not a stack is configured.
11. Click the **Add** button.
    The rule is added to the ACL.
12. Click the **Apply** button.
    Your settings are saved.
13. To save the settings to the running configuration, click the **Save** icon.

Modify an ACL rule that you created with the ACL Wizard

You can modify an ACL rule that you created with the ACL Wizard.

To modify an ACL rule that you created with the ACL Wizard:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > ACL Wizard**.
   The ACL Wizard page displays.
3. Select check box that is associated with the rule.
4. Update the match criteria as needed.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete an ACL rule that you created with the ACL Wizard

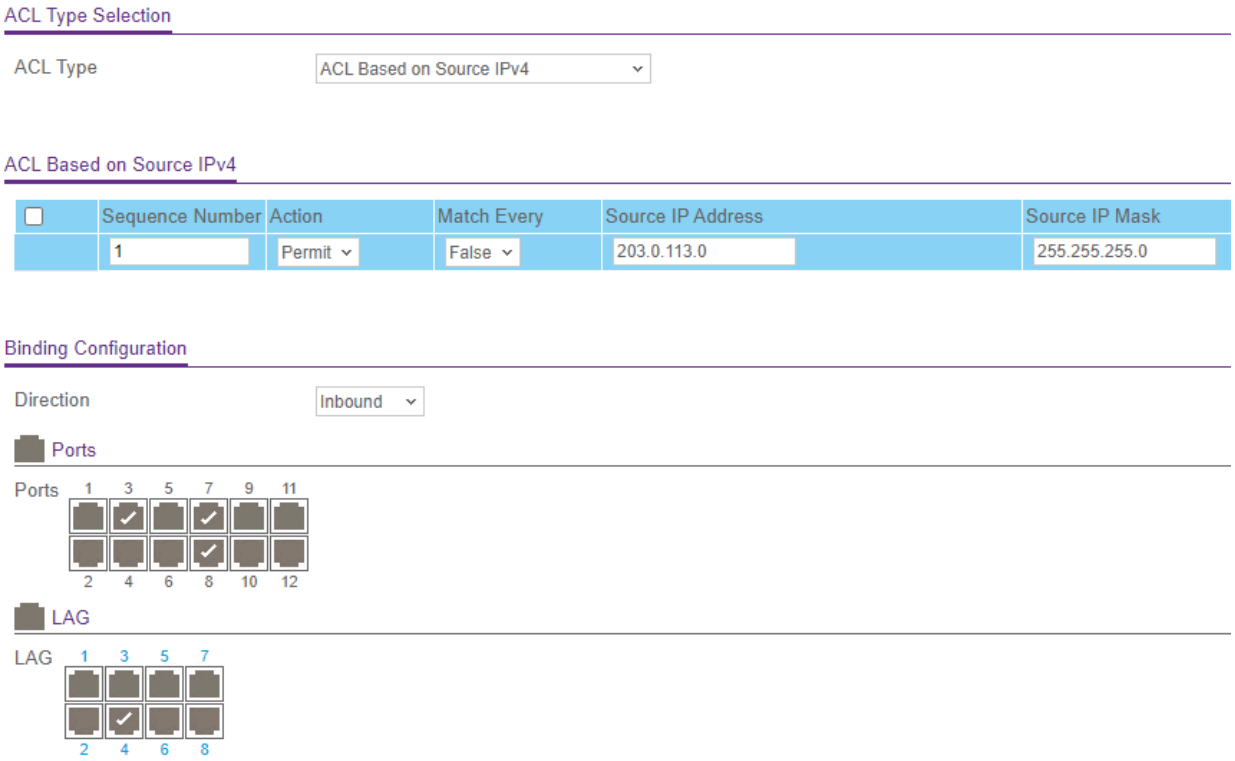
You can delete an ACL rule that you created with the ACL Wizard.

To delete an ACL rule that you created with the ACL Wizard:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > ACL Wizard**.
   The ACL Wizard page displays.
3. Select check box that is associated with the rule.
4. Click the **Delete** button.
   The rule is removed.
5. To save the settings to the running configuration, click the **Save** icon.

ACL Wizard example

In the following figure, the ACL rule is configured to check for packet matches on ports 3, 7, and 8 and on LAG 4. Packets that include a source address in the 203.0.113.0/24 network are permitted to be forwarded by the interfaces. All other packets are dropped because every ACL includes an implicit deny all rule as the last rule.



For information about the ACL Wizard, see Use the ACL Wizard to create a simple ACL on page 686.

Configure a MAC ACL

A MAC ACL consists of a set of rules that are matched sequentially against a packet. When a packet meets the match criteria of a rule, the specified rule action (Permit or Deny) is taken, and the additional rules are not checked for a match.
Multiple steps are involved in defining a MAC ACL and applying it to the switch:

1. Create a MAC ACL ID (see Add a MAC ACL on page 691).
2. Create a MAC rule (see Configure MAC ACL rules on page 693).
3. Associate the MAC ACL with one or more interfaces (see Configure MAC bindings on page 697).You can display or delete MAC ACL configurations in the MAC binding table (see Display or delete MAC ACL bindings in the MAC binding table on page 698).

Add a MAC ACL

You can add a MAC ACL, after which you can add a rule for the MAC ACL (see Add a rule for a MAC ACL on page 693).

To add a MAC ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select Security **Security > ACL > Basic > MAC ACL**.
   The MAC ACL page displays.
   The MAC ACL Table displays the number of ACLs currently configured on the switch and the maximum number of ACLs that can be configured. The current size is equal to the number of configured IPv4 and IPv6 ACLs plus the number of configured MAC ACLs.
3. In the **Name** field, specify a name for the MAC ACL.
   The name string can include alphabetic, numeric, hyphen, underscore, or space characters only. The name must start with an alphabetic character.
4. Click the **Add** button.
   The MAC ACL is added.
   Each configured ACL displays the following information:
   - **Rules**: The number of rules currently configured for the MAC ACL.
   - **Direction**: The direction of packet traffic affected by the MAC ACL, which can be Inbound or blank. (If the ACL is not bound to an interface, the direction is blank.)
5. To save the settings to the running configuration, click the **Save** icon.

Change the name of a MAC ACL

You can change the name of a MAC ACL.

To change the name of a MAC ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security** > **ACL > Basic > MAC ACL**.
   The MAC ACL page displays.
3. Select check box that is associated with the MAC ACL.
4. In the **Name** field, specify the new name.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete a MAC ACL

You can delete a MAC ACL that you no longer need.

To delete a MAC ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security** > **ACL > Basic > MAC ACL**.
   The MAC ACL page displays.
3. Select check box that is associated with the MAC ACL.
4. Click the **Delete** button.
   The MAC ACL is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure MAC ACL rules

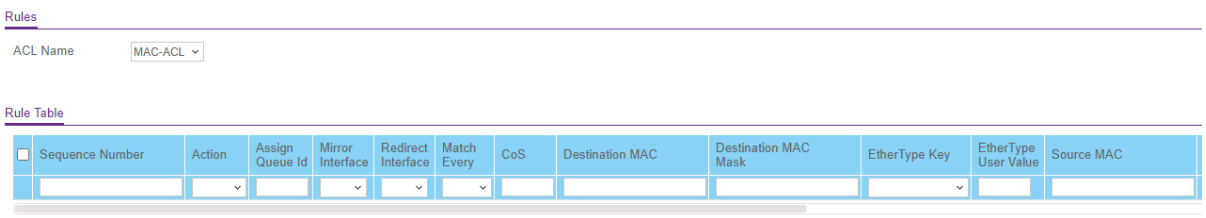
You can define rules for MAC-based ACLs. The access list definition includes rules that specify whether traffic matching the criteria is forwarded normally or discarded. A default deny all rule is the last rule of every list.

Add a rule for a MAC ACL

You can add a rule for a MAC ACL.

To add a rule for a MAC ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Basic > MAC Rules**.



3. From the **ACL Name** menu, select the MAC ACL.

4. In the **Sequence Number** field, enter a whole number in the range from 1 to 2147483647 to identify the rule.
5. From the **Action** menu, select the action that must be taken if a packet matches the rule's criteria:
   - **Permit**: Forwards packets that meet the ACL criteria.
   - **Deny**: Drops packets that meet the ACL criteria.
6. In the **Assign Queue ID** field, specify the hardware egress queue identifier that must be used to handle all packets matching this ACL rule.
   The range for the queue ID is from 0 to 6.
7. From the **Mirror Interface** menu, select the specific egress interface to which the matching traffic stream must be copied, in addition to being forwarded normally by the switch.
   This field cannot be set if a redirect interface is already configured for the ACL rule. This field is visible for a Permit action.
8. From the **Redirect Interface** menu, select the egress interface to which the matching traffic stream must be redirected, bypassing any forwarding decision normally performed by the switch.
   This field cannot be set if a mirror interface is already configured for the ACL rule.
9. From the **Match Every** menu, select whether each Layer 2 MAC packet must be matched against the rule:
   - **True**: Each packet must match the selected ACL rule.
   - **False**: Not all packets need to match the selected ACL rule.
10. In the **CoS** field, specify the 802.1p priority that must be compared against the information in an Ethernet frame.
    The range for the priority is from 0 to 7.
11. In the **Destination MAC** field, specify the destination MAC address that must be compared against the information in an Ethernet frame.
    The format is xx:xx:xx:xx:xx:xx. The BPDU keyword can be specified using a destination MAC address of 01:80:C2:xx:xx:xx.
12. In the **Destination MAC Mask** field, specify the destination MAC address mask that must be compared against the information in an Ethernet frame.
    The format is xx:xx:xx:xx:xx:xx. The BPDU keyword can be specified using a destination MAC mask of 00:00:00:ff:ff:ff.

13. From the **EtherType Key** menu, select the EtherType value that must be compared against the information in an Ethernet frame.
    The values are as follows:
    - **Apple Talk**
    - **IBM SNA**
    - **IPv4**
    - **IPv6**
    - **IPX**
    - **MPLS Multicast**
    - **MPLS Unicast**
    - **NetBios**
    - **Novell**
    - **PPPOE**
    - **RARP**
    - **User Value**
14. If you select **User Value** from the **EtherType Key** menu, in the **EtherType User Value** field, specify the customized EtherType value that must be used.
    This value must be compared against the information in an Ethernet frame. The range is from 0x0600 to 0xFFFF.
15. In the **Source MAC** field, specify the source MAC address that must be compared against the information in an Ethernet frame.
    The format is xx:xx:xx:xx:xx:xx.
16. In the **Source MAC Mask** field, specify the source MAC address mask that must be compared against the information in an Ethernet frame.
    The format is xx:xx:xx:xx:xx:xx.
17. In the **VLAN** field, specify the VLAN ID that must be compared against the information in an Ethernet frame.
    The range is from 1 to 4093. You can configure either a single VLAN or a VLAN range.
18. From the **Logging** menu, select whether to enable or disable logging.
    This option is available if the selection from the **Action** menu is **Deny**.
    If you select **Enable**, logging is enabled for this ACL rule (subject to resource availability on the switch).
    If the access list trap flag is also enabled, periodic traps are generated, indicating the number of times the rule was evoked during the report interval, which is fixed at five minutes.
19. In the **Rate Limit Conform Data Rate** field, specify the value of the conforming data rate, which is the data rate at which confirming traffic is limited.
    The range is from 1 to 4294967295 Kbps.

20. In the **Rate Limit Burst Size** field, specify the value of the burst size, which is the size at which bursts of traffic above the confirming data rate are permitted.
    The range is from 1 to 128 Kbps.
21. From the **Time Range** menu, as an option, select the name of the timer schedule.
    For information about timer schedules, see Timer schedules on page 153.
    The Rule Status field in the table shows if the rule is active or inactive. Blank means that no timer schedule is associated with the rule.
22. Click the **Add** button.
    The rule is added.
23. To save the settings to the running configuration, click the **Save** icon.

Change the match criteria for a MAC rule

You can change the match criteria for an existing MAC rule.

To change the match criteria for a MAC rule:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Basic > MAC Rules**.
   The MAC Rules page displays.
3. Select the check box that is associated with the rule.
4. Modify the fields as needed.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete a rule from a MAC ACL

You can delete a rule from a MAC ACL.

To delete a rule from a MAC:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Basic > MAC Rules**.
   The MAC Rules page displays.

3. Select the check box that is associated with the rule.
4. Click the **Delete** button.
   The rule is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure MAC bindings

When an ACL is bound to an interface, all the rules that are defined are applied to the selected interface. You can assign MAC ACLs to interfaces and LAGs.

To configure MAC binding:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Basic > MAC Binding Configuration**.
   The MAC Binding Configuration page displays.
3. From the **ACL ID** menu, select an ACL.
4. From the **Direction** menu, select the traffic direction for the MAC binding:
   - **Inbound**: The MAC binding can be applied to incoming traffic only.
   - **Outbound**: The MAC binding can be applied to outgoing traffic only.
5. In the **Sequence Number** field, optionally specify a number to indicate the order of the access list relative to other access lists already assigned to the interface and direction.
   A low number indicates high precedence order. If a sequence number is already in use for the interface and direction, the specified access list replaces the currently attached access list using that sequence number. If you do not specify the sequence number, a sequence number that is one number greater than the highest sequence number currently in use for the interface and direction is used. The range is from 1 to 4294967295.
6. To add the selected ACL to a port or LAG, in the Ports table (or if a stack is configured, in one of the Ports tables) or LAG table, click the port or LAG so that a check mark displays.
   The Ports and LAG tables display the available interfaces for ACL bindings. You can add the ACL to several ports and LAGs.
   The ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed. In addition, all LAGs are displayed, whether or not a stack is configured.
7. Click the **Apply** button.
   Your settings are saved.

8. To save the settings to the running configuration, click the **Save** icon.

   The following table describes the view-only fields on the page.

Table 191.Interface Binding Status information

Field	Description
Interface	The selected interface.
Direction	The selected packet filtering direction for the ACL.
ACL Type	The type of ACL assigned to the selected interface and direction.
ACL ID	The ACL number or name identifying the ACL assigned to the selected interface and direction.
Sequence Number	The sequence number signifying the order of specified ACL relative to other ACLs assigned to the selected interface and direction.

Display or delete MAC ACL bindings in the MAC binding table

You can display or delete the MAC ACL bindings in the MAC binding table.

To display or delete MAC ACL bindings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.

   The System Information page displays.
2. Select **Security > ACL > Basic > Binding Table**.

   The Binding Table page displays. This table displays the MAC ACL-to-interface bindings.
3. To delete a MAC ACL-to-interface binding, do the following:
   a. Select the check box next to the interface.
   b. Click the **Delete** button.

      The binding is removed.
4. To save the settings to the running configuration, click the **Save** icon.

The following table describes the information that is displayed in the MAC binding table.

Table 192.MAC Binding Table information

Field	Description
Interface	The interface to which the ACL is bound.
Direction	The selected packet filtering direction for the ACL.
ACL Type	The type of ACL assigned to selected interface and direction.
ACL ID	The ACL name identifying the ACL assigned to selected interface and direction.
Sequence Number	The sequence number signifying the order of the specified ACL relative to other ACLs assigned to selected interface and direction.

Configure a basic or extended IPv4 ACL

An IPv4 ACL consists of a set of rules that are matched sequentially against a packet. When a packet meets the match criteria of a rule, the specified rule action (Permit or Deny) is taken, and the additional rules are not checked for a match. You must specify the interfaces to which an IPv4 ACL applies, as well as whether it applies to inbound or outbound traffic.
Multiple steps are involved in defining an IPv4 ACL and applying it to the switch:

1. Add an IPv4 ACL ID (see Add an IPv4 ACL on page 700).
   The differences between a basic IPv4 ACL and an extended IPv4 ACL are as follows:
   - **Numbered ACL from 1 to 99**: Creates a basic IPv4 ACL, which allows you to permit or deny traffic from a source IP address.
   - **Numbered ACL from 100 to 199**: Creates an extended IPv4 ACL, which allows you to permit or deny specific types of Layer 3 or Layer 4 traffic from a source IP address to a destination IP address. This type of ACL provides more granularity and filtering capabilities than the basic IP ACL.
   - **Named IP ACL**: Creates an extended IPv4 ACL with a name string that is up to 31 alphanumeric characters in length. The name must start with an alphabetic character.
2. Create an IPv4 rule (see Configure rules for a basic IP ACL on page 701 or Configure rules for an extended IPv4 ACL on page 705).
3. Associate the IPv4 ACL with one or more interfaces (see Configure IP ACL interface bindings on page 720).
   You can display or delete IPv4 ACL configurations in the IP ACL Binding table (see Display or delete IP ACL bindings in the IP ACL binding table on page 722.

Add an IPv4 ACL

You can add an IPv4 ACL, after which you can add a rule for the IPv4 ACL (see Add a rule for a basic IPv4 ACL on page 702 or Add a rule for an extended IPv4 ACL on page 705).

To add an IPv4 ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IP ACL**.
   The IP ACL page displays.
   The IP ACL page shows the current size of the ACL table compared to the maximum size of the ACL table. The current size is equal to the number of configured IPv4 ACLs plus the number of configured MAC ACLs and IPv6 ACLs. The maximum number of ACLs on the switch is 100.
   The Current Number of ACL field displays the current number of all ACLs configured on the switch.
   The Maximum ACL field displays the maximum number of ACLs that you can configure on the switch.
3. In the **IP ACL ID** field, specify the ACL ID or IP ACL name, which depends on the IP ACL type.
   The IP ACL ID is an integer in the following range:
   - **1–99**: Creates a basic IP ACL, which allows you to permit or deny traffic from a source IP address.
   - **100–199**: Creates an extended IP ACL, which allows you to permit or deny specific types of Layer 3 or Layer 4 traffic from a source IP address to a destination IP address. This type of ACL provides more granularity and filtering capabilities than the standard IP ACL.
   - **IP ACL Name**: Creates an extended IP ACL with a name string that is up to 31 alphanumeric characters in length. The name must start with an alphabetic character.

   Each configured ACL displays the following information:
   - **Rules**: The number of rules currently configured for the IPv4 ACL.
   - **Type**: Identifies the ACL as a basic IP ACL (with ID from 1 to 99), extended IP ACL (with ID from 100 to 199), or named ACL (which is also an extended ACL).
4. Click the **Add** button.
   The IP ACL is added.
5. To save the settings to the running configuration, click the **Save** icon.

Change the number or name of an IPv4 ACL

You can change the number or name of an IPv4 ACL.

To change the number or name of an IPv4 ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IP ACL**.
   The IP ACL Configuration page displays.
3. Select the check box that is associated with the IP ACL.
4. In the **IP ACL** field, specify the new number or name.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete an IPv4 ACL

You can delete an IPv4 ACL that you no longer need.

To delete an IPv4 ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IP ACL**.
   The IP ACL Configuration page displays.
3. Select the check box that is associated with the IP ACL.
4. Click the **Delete** button.
   The IP ACL is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure rules for a basic IP ACL

You can define rules for IP-based standard ACLs (basic ACLs). The access list definition includes rules that specify whether traffic matching the criteria is forwarded normally or discarded.

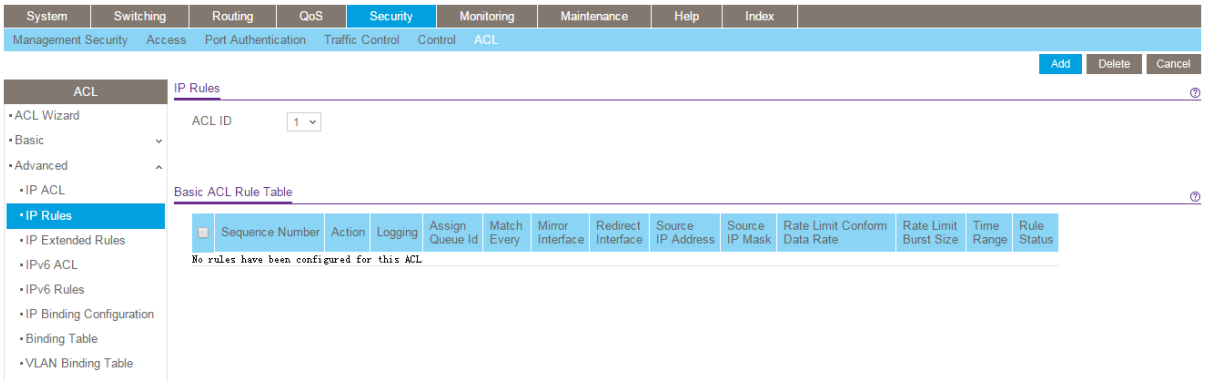
NOTE: An implicit deny all rule is included at the end of an ACL list. This means that if an ACL is applied to a packet, and if none of the explicit rules match, then the final implicit deny all rule applies and the packet is dropped.

Add a rule for a basic IPv4 ACL

You can add a rule for a basic IPv4 ACL.

To add a rule for a basic IPv4 ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IP Rules**.



   If no rules exists, the Basic ACL Rule Table shows the message *No rules have been configured for this ACL*. If one or more rule exists for the ACL, the rules display in the Basic ACL Rule Table.
3. From the **ACL ID** menu, select the IP ACL for which you want to add or a change a rule.
   For basic IP ACLs, this must be an ID in the range from 1 to 99.
4. Click the **Add** button.
   The page that displays lets you configure the basic (standard) ACL rule.

5. Specify the following match criteria for the rule:
   - **Sequence Number**: Enter an ACL sequence number in the range from 1 to 2147483647 that is used to identify the rule. An IP ACL can contain up to 50 rules.
   - **Action**: Select the ACL forwarding action, which is one of the following:
     - **Permit**: Forward packets that meet the ACL criteria. **Egress Queue**: If the selection from the **Action** menu is **Permit**, you can specify the hardware egress queue identifier that is used to handle all packets matching this IP ACL rule. The range of queue IDs is from 0 to 6.
     - **Deny**: Drop packets that meet the ACL criteria. **Logging**: If the selection from the **Action** menu is **Deny**, you can enable logging for the ACL by selecting the **Enable** radio button. (Logging is subject to resource availability on the switch.)
   - **Match Every**: Select one of the radio buttons to specify whether all packets must match the selected IP ACL rule:
     - **Enable**: All packets must match the selected IP ACL rule and are either permitted or denied.
     - **Disable**: Not all packets need to match the selected IP ACL rule.
   - **Mirror Interface**: From the menu, select the egress interface to which the matching traffic stream must be copied, in addition to being forwarded normally by the switch. You can either select a mirror interface or a redirect interface. These options are mutually exclusive.
   - **Redirect Interface**: From the menu, select the egress interface to which the matching traffic stream must be redirected, bypassing any forwarding decision normally performed by the switch. You can either select a redirect interface or a mirror interface. These options are mutually exclusive.
   - **Src IP Address**: Specify an IP address using dotted-decimal notation to be compared to a packet's source IP address as a match criterion for the selected IP ACL rule.
   - **Src IP Mask**: Specify the IP mask in dotted-decimal notation to be used with the source IP address.
   - **Rate Limit Conform Data Rate**: Specify the value of the conforming data rate, which is the data rate at which confirming traffic is limited. The range is from 1 to 4294967295 Kbps.
   - **Rate Limit Burst Size**: Specify the value of the burst size, which is the size at which bursts of traffic above the confirming data rate are permitted. The range is from 1 to 128 Kbps.
   - **Time Range**: From the **Time Range** menu, select the timer schedule that must be associated with the rule. For more information about timer schedules, see Timer schedules on page 153. If you did not set up any timer schedules, the menu does not present any options.

6. Click the **Apply** button.
   Your settings are saved.
   The Rule Status field in the table shows if the rule is active or inactive. Blank means that no timer schedule is associated with the rule.
7. To save the settings to the running configuration, click the **Save** icon.

Modify the match criteria for a basic IPv4 ACL rule

You can modify the match criteria for a basic IPv4 ACL rule.

To modify the match criteria for a basic IPv4 ACL rule:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IP Rules**.
   The IP Rules page displays.
3. From the **ACL ID** menu, select the ACL that includes the rule that you want to modify.
4. In the Basic ACL Rule Table, click the rule.
   The rule is a hyperlink. The Standard ACL Rule Configuration page displays.
5. Modify the basic IP ACL rule criteria.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Delete a basic IPv4 ACL rule

You can delete a basic IPv4 ACL rule that you no longer need.

To delete a basic IPv4 ACL rule:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IP Rules**.
   The IP Rules page displays.
3. From the **ACL ID** menu, select the ACL that includes the rule that you want to modify.
4. In the Basic ACL Rule Table, select the check box that is associated with the rule.

5. Click the **Delete** button.
   The rule is removed.
6. To save the settings to the running configuration, click the **Save** icon.

Configure rules for an extended IPv4 ACL

You can define rules for extended IPv4 ACLs. The access list definition includes rules that specify whether traffic matching the criteria is forwarded normally or discarded.

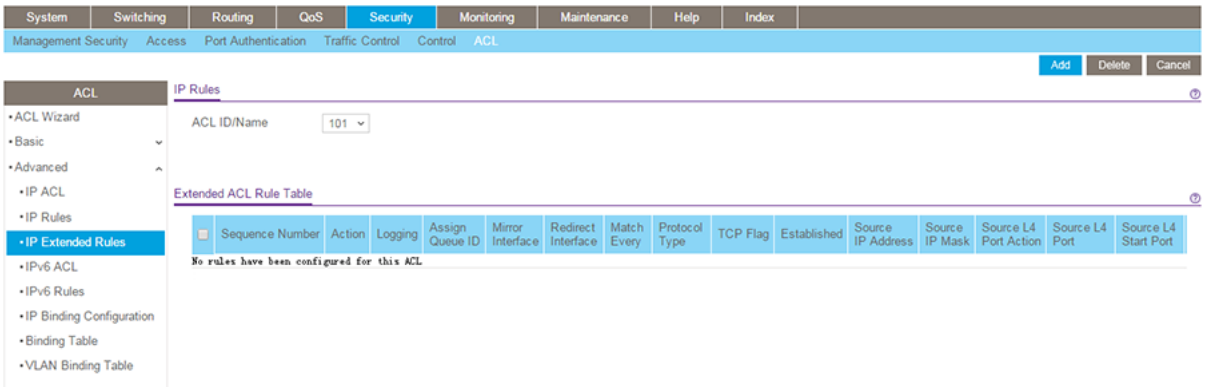
NOTE: An implicit deny all rule is included at the end of an ACL list. This means that if an ACL is applied to a packet and if none of the explicit rules match, then the final implicit deny all rule applies and the packet is dropped.

Add a rule for an extended IPv4 ACL

You can add a rule for an extended IPv4 ACL.

To add a rule for an extended IPv4 ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IP Extended Rules**.



If no rules exists, the Extended ACL Rule Table shows the message *No rules have been configured for this ACL*. If one or more rule exists for the ACL, the rules display in the Extended ACL Rule Table.

3. From the **ACL ID/Name** menu, select the IP ACL for which you want to add a rule.
   For extended IP ACLs, this must be an ID in the range from 101 to 199 or a name.

4. Click the **Add** button.
   The page that displays lets you configure the extended rule.

5. Configure the following options for the rule:
   - **Sequence Number**: Enter a number in the range from 1 to 2147483647 that is used to identify the rule. An extended IP ACL can contain up to 1023 rules.
   - **Action**: Select the ACL forwarding action, which is one of the following:
     - **Permit**: Forward packets that meet the ACL criteria. **Egress Queue**: If the selection from the **Action** menu is **Permit**, select the hardware egress queue identifier that is used to handle all packets matching this IP ACL rule. The range of queue IDs is from 0 to 6.
     - **Deny**: Drop packets that meet the ACL criteria. This option is available if the selection from the **Action** menu is **Deny**.
       If you select **Enable**, logging is enabled for this ACL rule (subject to resource availability on the switch).
       If the access list trap flag is also enabled, periodic traps are generated, indicating the number of times the rule was evoked during the report interval, which is fixed at five minutes.
   - **Interface**: For a Permit action, use either a mirror interface or a redirect interface:
     - Select the **Mirror** radio button and use the menu to specify the egress interface to which the matching traffic stream is copied, in addition to being forwarded normally by the device.
     - Select the **Redirect** radio button and use the menu to specify the egress interface to which the matching traffic stream is forced, bypassing any forwarding decision normally performed by the device.
   - **Match Every**: Select one of the radio buttons to specify whether all packets must match the selected IP ACL rule:
     - **False**: Not all packets need to match the selected IP ACL rule. You can configure other match criteria on the page.
     - **True**: All packets must match the selected IP ACL rule and are either permitted or denied. In this case, you cannot configure other match criteria on the page.
   - **Protocol Type**:From the menu, select a protocol that a packet's IP protocol must be matched against: **IP**, **ICMP**, **IGMP**, **TCP**, **UDP**, **EIGRP**, **GRE**, **IPINIP**, **OSPF**, **PIM**, or **Other**. If you select **Other**, enter a protocol number from 0 to 255.
   - **TCP Flag**: If you select **TCP** from the **Protocol Type** menu, for each TCP flag, you can specify whether or not a packet's TCP flag must match. The TCP flag values are URG, ACK, PSH, RST, SYN, and FIN. You can set each TCP flag separately to one of the following options:
     - **Ignore**: The packet's TCP flag is ignored. This is the default setting.
     - **Set**: A packet matches this ACL rule if the TCP flag in this packet is set.
     - **Clear**: A packet matches this ACL rule if the TCP flag in this packet is not set.

> **NOTE:** If the RST and ACK flags are set, the option **Established** is available, indicating that a match occurs if either the RST- or ACK-specified bits are set in the packet's header.

- **Src**: In the **Src** field, enter a source IP address, using dotted-decimal notation, to be compared to a packet's source IP address as a match criterion for the selected IP ACL rule:
  - If you select the **IP Address** radio button, enter an IP address or an IP address range. You can enter a relevant wildcard mask to apply this criteria. If this field is left empty, it means any.
  - If you select the **Host** radio button, the wildcard mask is configured as 0.0.0.0. If this field is left empty, it means any. The wildcard mask determines which bits are used and which bits are ignored. A wildcard mask of 0.0.0.0 indicates that none of the bits are important. A wildcard of 255.255.255.255 indicates that all of the bits are important.
- **Src L4**: The options are available only if the selection from the **Protocol Type** menu is **TCP** or **UDP**. Use the source L4 port option to specify relevant matching conditions for L4 port numbers in the extended ACL rule. You can select either the **Port** radio button or the **Range** radio button:
  - **Port**: If you select the Port radio button, you can either enter the port number yourself or select one of the following protocols from the menu:
    - The source IP TCP port protocols are **Domain**, **Echo**, **FTP**, **FTP data**, **www-http**, **SMTP**, **Telnet**, **POP2**, **POP3**, and **BGP**.
    - The source IP UDP port protocols are **Domain**, **Echo**, **SNMP**, **NTP**, **RIP**, **Time**, **Who**, and **TFTP**.

    Each of these values translates into its equivalent port number, which is used as both the start and end of the port range.

    Select Other from the menu to enter a port number. If you select Other from the menu but leave the field blank, it means any.

    The relevant matching conditions for L4 port numbers are as follows:
    - **Equal**: The IP ACL rule matches if the Layer 4 source port number is equal to the specified port number or port key.
    - **Not Equal**: The IP ACL rule matches if the Layer 4 source port number is not equal to the specified port number or port key.
    - **Less Than**: The IP ACL rule matches if the Layer 4 source port number is a lower number than the specified port number or port key.
    - **Greater Than**: The IP ACL rule matches if the Layer 4 source port number is a higher number than the specified port number or port key.
  - **Range**: If you select the Range radio button, the IP ACL rule matches only if the

Layer 4 source port number is within the specified port range. The starting port, ending port, and all ports in between are a part of the Layer 4 port range. The **Start Port** and **End Port** fields identify the first and last ports that are part of the port range. The values can range from 0 to 65535.

You can either enter the port range yourself or select one of the following protocols from the menu:

- The source IP TCP port protocols are **Domain**, **Echo**, **FTP**, **FTP data**, **www-http**, **SMTP**, **Telnet**, **POP2**, **POP3**, and **BGP**.
- The source IP UDP port protocols are **Domain**, **Echo**, **SNMP**, **NTP**, **RIP**, **Time**, **Who**, and **TFTP**.

Each of these values translates into its equivalent port number, which is used as both the start and end of the port range. Select Other from the menu to enter a port number. If you select Other from the menu but leave the field blank, it means any.

The wildcard mask determines which bits are used and which bits are ignored. A wildcard mask of 0.0.0.0 indicates that none of the bits are important. A wildcard of 255.255.255.255 indicates that all of the bits are important.

- **Dst**: In the **Dst** field, enter a destination IP address, using dotted-decimal notation, to be compared to a packet's destination IP address as a match criteria for the selected IP ACL rule:
  - If you select the **IP Address** radio button, enter an IP address with a relevant wildcard mask to apply this criteria. If this field is left empty, it means any.
  - If you select the **Host** radio button, the wildcard mask is configured as 0.0.0.0. If this field is left empty, it means any. The wildcard mask determines which bits are used and which bits are ignored. A wildcard mask of 0.0.0.0 indicates that none of the bits are important. A wildcard of 255.255.255.255 indicates that all of the bits are important.
- **Dst L4**: The options are available only if the selection from the **Protocol Type** menu is **TCP** or **UDP**. Use the destination L4 port option to specify relevant matching conditions for L4 port numbers in the extended ACL rule.

  You can select either the **Port** radio button or the **Range** radio button:
  - **Port**: If you select the Port radio button, you can either enter the port number yourself or select one of the following protocols from the menu.
    - The destination IP TCP port protocols are **Domain**, **Echo**, **FTP**, **FTP data**, **www-http**, **SMTP**, **Telnet**, **POP2**, **POP3**, and **BGP**.
    - The destination IP UDP port protocols are **Domain**, **Echo**, **SNMP**, **NTP**, **RIP**, **Time**, **Who**, and **TFTP**.

    Each of these values translates into its equivalent port number, which is used as both the start and end of the port range.

    Select Other from the menu to enter a port number. If you select Other from the menu but leave the field blank, it means any.

The relevant matching conditions for L4 port numbers are as follows:

- **Equal**: The IP ACL rule matches if the Layer 4 source port number is equal to the specified port number or port key.
- **Not Equal**: The IP ACL rule matches if the Layer 4 source port number is not equal to the specified port number or port key.
- **Less Than**: The IP ACL rule matches if the Layer 4 source port number is a lower number than the specified port number or port key.
- **Greater Than**: The IP ACL rule matches if the Layer 4 source port number is a higher number than the specified port number or port key.

- **Range**: If you select the Range radio button, the IP ACL rule matches only if the Layer 4 destination port number is within the specified port range. The starting port, ending port, and all ports in between are a part of the Layer 4 port range. The **Start Port** and **End Port** fields identify the first and last ports that are part of the port range. They values can range from 0 to 65535.
  You can either select the enter the port range yourself or select one of the following protocols from the menu:
  - The destination IP TCP port protocols are **Domain**, **Echo**, **FTP**, **FTP data**, **www-http**, **SMTP**, **Telnet**, **POP2**, **POP3**, and **BGP**.
  - The destination IP UDP port protocols are **Domain**, **Echo**, **SNMP**, **NTP**, **RIP**, **Time**, **Who**, and **TFTP**.

  Each of these values translates into its equivalent port number, which is used as both the start and end of the port range.
  Select Other from the menu to enter a port number. If you select Other from the menu but leave the field blank, it means any.
  The wildcard mask determines which bits are used and which bits are ignored. A wildcard mask of 0.0.0.0 indicates that none of the bits are important. A wildcard of 255.255.255.255 indicates that all of the bits are important.

- **IGMP Type**: If your selection from the **Protocol Type** menu is **IGMP** and you specify the IGMP type, the IP ACL rule matches the specified IGMP message type. The range is from 0 to 255. If this field is left empty, it means any.
- **ICMP**: If your selection from the **Protocol Type** menu is **ICMP**, you can select either the **Type** or **Message** radio button:
  - If you select the **Type** radio button, note the following:
    - The **Type** and **Code** fields are enabled only if the protocol is ICMP. Use these fields to specify a match condition for ICMP packets:
    - If you specify information in the **Type** field, the IP ACL rule matches the specified ICMP message type. The type number can be from 0 to 255.
    - If you specify information in the **Code** field, the IP ACL rule matches the specified ICMP message code. The code can be from 0 to 255.
    - If these fields are left empty, it means any.

        - If you select the **Message** radio button, from the menu, select the type of the ICMP message to match with the selected IP ACL rule. Specifying a type of message implies that both the ICMP type and ICMP code are specified. The ICMP message is decoded into the corresponding ICMP type and ICMP code within the ICMP type. The IPv4 ICMP message types are **Echo**, **echo-reply**, **host-redirect**, **mobile-redirect**, **net-redirect**, **net-unreachable**, **redirect**, **packet-too-big**, **port-unreachable**, **source-quench**, **router-solicitation**, **router-advertisement**, **TTL-exceeded**, **time-exceeded**, and **unreachable**.
    - **Fragments**: Either select the **Enable** radio button to allow initial fragments (that is, the fragment bit is asserted) or leave the default **Disable** radio button selected to prevent initial fragments from being used. This option is not valid for rules that match L4 information such as a TCP port number, because that information is carried in the initial packet.
    - **Service Type**: Select a service type match condition for the extended IP ACL rule. The possible values are **IP DSCP**, **IP precedence**, and **IP TOS**, which are alternative methods to specify a match criterion for the same service type field in the IP header. Each method uses a different user notation. After you make a selection is made, you can specify the appropriate values.
        - **IP DSCP**: This is an optional configuration. Specify the IP DiffServ Code Point (DSCP) field. The DSCP is defined as the high-order 6 bits of the service type octet in the IP header. Enter an integer from 0 to 63. To select the IP DSCP, select one of the DSCP keywords from the menu. To specify a numeric value, select Other and a field displays in which you can enter numeric value of the DSCP.
        - **IP Precedence**: This is an optional configuration. The IP precedence field in a packet is defined as the high-order three bits of the service type octet in the IP header. Enter an integer from 0 to 7.
        - **IP TOS**: This is an optional configuration. The IP ToS field in a packet is defined as all 8 bits of the service type octet in the IP header. The ToS bits value is a hexadecimal number from 00 to 09 and to aa to ff. The ToS mask value is a hexadecimal number from 00 to FF. The ToS mask denotes the bit positions in the ToS bits value that are used for comparison against the IP ToS field in a packet. For example, to check for an IP ToS value for which bit 7 is set and is the most significant value, for which bit 5 is set, and for which bit 1 is cleared, use a ToS bits value of 0xA0 and a ToS mask of 0xFF.
    - **Rate Limit Conform Data Rate**: Specify the value of the conforming data rate, which is the data rate at which confirming traffic is limited. The range is from 1 to 4294967295 Kbps.
    - **Rate Limit Burst Size**: Specify the value of the burst size, which is the size at which bursts of traffic above the confirming data rate are permitted. The range is from 1 to 128 Kbps.
    - **Time Range**: From the **Time Range** menu, select the timer schedule that must be

associated with the rule. For more information about timer schedules, see Timer schedules on page 153. If you did not set up any timer schedules, the menu does not present any options.

6. Click the **Apply** button.
   Your settings are saved.
   The Rule Status field in the table shows if the rule is active or inactive. Blank means that no timer schedule is associated with the rule.
7. To save the settings to the running configuration, click the **Save** icon.

Modify the match criteria for an extended IPv4 ACL rule

You can modify the match criteria for an extended IPv4 ACL rule.

To modify the match criteria for an existing extended IPv4 ACL rule:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IP Extended Rules**.
   The IP Extended Rules page displays.
3. From the **ACL ID/Name** menu, select the ACL that includes the rule that you want to modify.
4. In the Extended ACL Rule Table, click the rule.
   The rule is a hyperlink. The Extended ACL Rule Configuration page displays.
5. Modify the extended IP ACL rule criteria.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Delete an extended IPv4 ACL rule

You can delete an extended IPv4 ACL rule that you no longer need.

To delete an extended IPv4 ACL rule:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IP Extended Rules**.
   The IP Extended Rules page displays.

3. From the **ACL ID/Name** menu, select the ACL that includes the rule that you want to delete.
4. In the Extended ACL Rule Table, select the check box that is associated with the rule.
5. Click the **Delete** button.
   The rule is removed.
6. To save the settings to the running configuration, click the **Save** icon.

Configure an IPv6 ACL

An IPv6 ACL consists of a set of rules that are matched sequentially against a packet. When a packet meets the match criteria of a rule, the specified rule action (Permit or Deny) is taken, and the additional rules are not checked for a match. You must specify the interfaces to which an IPv6 ACL applies, as well as whether it applies to inbound or outbound traffic.
Multiple steps are involved in defining an IPv6 ACL and applying it to the switch:

1. Add an IPv6 ACL ID (see Add an IPv6 ACL on page 713).An IPv6 ACL must start with a name string that is up to 31 alphanumeric characters in length. The name must start with an alphabetic character.
2. Create an IPv6 rule (see Configure rules for an IPv6 ACL on page 715).
3. Associate the IPv6 ACL with one or more interfaces (see Configure IP ACL interface bindings on page 720).You can display or delete IPv6 ACL configurations in the IP ACL Binding table (see Display or delete IP ACL bindings in the IP ACL binding table on page 722).

Add an IPv6 ACL

You can add an IPv6 ACL, after which you can add a rule for the IPv6 ACL (see Add a rule for an IPv6 ACL on page 715).

To add an IPv6 ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IPv6 ACL**.
   The IPv6 ACL page displays.
3. In the **IPv6 ACL** field, specify a name to identify the IPv6 ACL.
   This is the IPv6 ACL name string, which includes up to 31 alphanumeric characters only. The name must start with an alphabetic character.
4. Click the **Add** button.
   The IPv6 ACL is added to the switch configuration.

5. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 193.IPv6 ACL information

Field	Description
Current Number of ACL	The current number of ACLs configured on the switch.
Maximum ACL	The maximum number of IP ACLs that can be configured on the switch.
Rules	The maximum number of ACLs that can be configured on the switch.
Type	The type is IPv6 ACL.

Change the name of an IPv6 ACL

You can change the name of an IPv6 ACL.

To change the name of an IPv6 ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IPv6 ACL**.
   The IPv6 ACL page displays.
3. Select the check box that is associated with the IPv6 ACL.
4. In the **IPv6 ACL** field, specify the new name.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Delete an IPv6 ACL

You can delete an IPv6 ACL that you no longer need.

To delete an IPv6 ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Security > ACL > Advanced > IPv6 ACL**.
   The IPv6 ACL page displays.
3. Select the check box that is associated with the IPv6 ACL.
4. Click the **Delete** button.
   The IPv6 ACL is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Configure rules for an IPv6 ACL

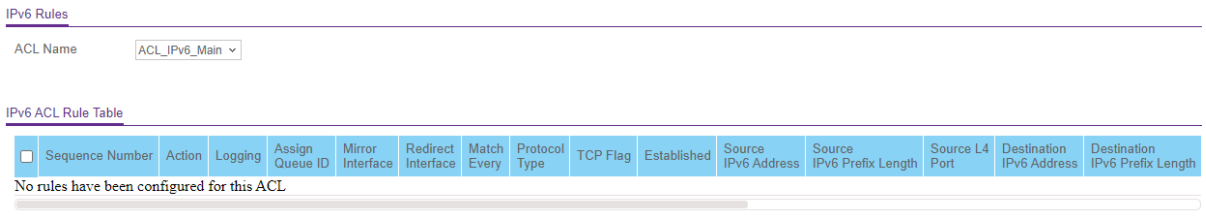
You can define rules for IPv6 ACLs. The access list definition includes rules that specify whether traffic matching the criteria is forwarded normally or discarded.

Add a rule for an IPv6 ACL

You can add a rule for an IPv6 ACL.

To add a rule for an IPv6 ACL:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IPv6 Rules**.



   If no rules exists, the IPv6 ACL Rule Table shows the message *No rules have been configured for this ACL*. If one or more rule exists for the ACL, the rules display in the IPv6 ACL Rule Table.
3. From the **ACL Name** menu, select the IPv6 ACL for which you want to add or a change a rule.
4. Click the **Add** button.
   The page that displays lets you configure the IPv6 rule.

5. Configure the following options for the rule:
   - **Sequence Number**: Enter a whole number in the range of 1 to 2147483647. This number is used to identify the rule. An IPv6 ACL can contain up to 1023 rules.
   - **Action**: Select the ACL forwarding action by selecting one of the following radio buttons:
     - **Permit**: Forward packets that meet the ACL criteria.
     - **Deny**: Drop packets that meet the ACL criteria.
   - **Egress Queue**: If you select the **Permit** radio button, select the hardware egress queue identifier that is used to handle all packets matching this IPv6 ACL rule. The range of queue IDs is from 0 to 6.
   - **Logging**: If you select the **Deny** radio button, you can enable logging for the ACL by selecting the **Enable** radio button. (Logging is subject to resource availability in the device.) If the access list trap flag is also enabled, periodic traps are generated, indicating the number of times this rule was evoked during the report interval. A fixed five-minute report interval is used for the switch. A trap is not issued if the ACL rule hit count is zero for the current interval.
   - **Interface**: If you select the **Permit** radio button, use either a mirror interface or a redirect interface:
     - Select the **Mirror Interface** radio button and use the menu to specify the egress interface to which the matching traffic stream is copied, in addition to being forwarded normally by the device.
     - Select the **Redirect Interface** radio button and use the menu to specify the egress interface to which the matching traffic stream is forced, bypassing any forwarding decision normally performed by the device.
   - **Match Every**: Select whether all packet must match the selected IPv6 ACL rule:
     - **Disable**: Not all packets need to match the selected IPv6 ACL rule. You can configure other match criteria on the page.
     - **Enable**: All packets must match the selected IPv6 ACL rule and are either permitted or denied. In this case, you cannot configure other match criteria on the page.
   - **Protocol Type**: Specify the IPv6 protocol type in one of the following ways:
     - From the **Protocol Type** menu, select **IPv6**, **ICMPv6**, **TCP**, or **UDP**.
     - From the **Protocol Type** menu, select **Other**, and in the associated field, specify an integer ranging from 0 to 255. This number represents the IPv6 protocol.
   - **TCP Flag**: If you select **TCP** from the **Protocol Type** menu, for each TCP flag, you can specify whether or not a packet's TCP flag must match. The TCP flag values are URG, ACK, PSH, RST, SYN, and FIN. You can set each TCP flag separately to one of the following options:
     - **Ignore**: The packet's TCP flag is ignored. This is the default setting.
     - **Set**: A packet matches this ACL rule if the TCP flag in this packet is set.
     - **Clear**: A packet matches this ACL rule if the TCP flag in this packet is not set.

NOTE: If the RST and ACK flags are set, the option **Established** is available, indicating that a match occurs if either the RST- or ACK-specified bits are set in the packet's header.

- **Src**: In the **Src** field, enter a source IPv6 address or source IPv6 address range to be compared to a packet's source IPv6 address as a match criterion for the selected IPv6 ACL rule:
  - If you select the **IPv6 Address** radio button, enter an IPv6 address or IPv6 range to apply this criteria. If this field is left empty, it means any.
  - If you select the **Host** radio button, enter a host source IPv6 address to match the specified IPv6 address. If this field is left empty, it means any.

  The source IPv6 address argument must be in the form documented in RFC 2373 where the address is specified in hexadecimal numbers using 16-bit values between colons.
- **Src L4**: The options are available only if the selection from the **Protocol Type** menu is **TCP** or **UDP**. Use the source L4 port option to specify relevant matching conditions for L4 port numbers in the IPv6 ACL rule.

  You can either enter the port number yourself or select one of the following protocols from the menu:
  - The source IP TCP port protocols are **Domain**, **Echo**, **FTP**, **FTP data**, **www-http**, **SMTP**, **Telnet**, **POP2**, **POP3**, and **BGP**.
  - The source IP UDP port protocols are **Domain**, **Echo**, **SNMP**, **NTP**, **RIP**, **Time**, **Who**, and **TFTP**.

  Each of these values translates into its equivalent port number, which is used as both the start and end of the port range.

  Select **Other** from the menu to enter a port number. If you select Other from the menu but leave the field blank, it means any.

  The only relevant matching condition for L4 port numbers is equal. This means that an IPv6 ACL rule matches only if the Layer 4 source port number is equal to the specified port number or port protocol.
- **Dst**: In the **Dst** field, enter a destination IPv6 address to be compared to a packet's destination IPv6 address as a match criterion for the selected IPv6 ACL rule:
  - If you select the **IPv6 Address** radio button, enter an IPv6 address to apply this criteria. If this field is left empty, it means any.
  - If you select the **Host** radio button, enter a host source IPv6 address to match the specified IPv6 address. If this field is left empty, it means any.

  The source IPv6 address argument must be in the form documented in RFC 2373 where the address is specified in hexadecimal numbers using 16-bit values between colons.
- **Dst L4**: The options are available only if the selection from the **Protocol Type** menu is **TCP** or **UDP**. Use the destination L4 port option to specify relevant matching conditions

for L4 port numbers in the IPv6 ACL rule.
You can either enter the port number yourself or select one of the following protocols from the menu:

- The destination IP TCP port protocols are **Domain**, **Echo**, **FTP**, **FTP data**, **www-http**, **SMTP**, **Telnet**, **POP2**, **POP3**, and **BGP**.
- The destination IP UDP port protocols are **Domain**, **Echo**, **SNMP**, **NTP**, **RIP**, **Time**, **Who**, and **TFTP**.

Each of these values translates into its equivalent port number, which is used as both the start and end of the port range.
Select **Other** from the menu to enter a port number. If you select Other from the menu but leave the field blank, it means any.
The only relevant matching condition for L4 port numbers is equal. This means that an IPv6 ACL rule matches only if the Layer 4 destination port number is equal to the specified port number or port protocol.

- **ICMPv6**: The **Type** and **Message** fields are enabled only if the protocol is ICMPv6. Use these fields to specify a match condition for ICMPv6 packets. Select either the **Type** or **Message** radio button:
  - If you select the **Type** radio button, note the following:
    - The **Type** and **Code** fields are enabled only if the protocol is ICMPv6. Use these fields to specify a match condition for ICMPv6 packets:
    - If you specify information in the **Type** field, the IPv6 ACL rule matches the specified ICMPv6 message type. The type number can be from 0 to 255.
    - If you specify information in the **Code** field, the IPv6 ACL rule matches the specified ICMPv6 message code. The code can be from 0 to 255.
    - If these fields are left empty, it means any.
  - If you select the **Message** radio button, note the following:
    - Select the type of the ICMPv6 message to match with the selected IPv6 ACL rule. Specifying a type of message implies that both the ICMPv6 type and ICMPv6 code are specified. The ICMPv6 message is decoded into the corresponding ICMPv6 type and ICMPv6 code within the ICMP type.
    - The ICMPv6 message types are **destination-unreachable**, **echo-reply**, **echo-request**, **header**, **hop-limit**, **MLD-query**, **MLD-reduction**, **MLD-report**, **next-header**, **no-admin**, **no-route**, **packet-too-big**, **port-unreachable**, **router-solicitation**, **router-advertisement**, **router-renumbering**, **unreachable**, **time-exceeded**, **nd-na**, and **nd-ns**.
- **Fragments**: Either select the **Enable** radio button to allow initial fragments (that is, the fragment bit is asserted) or leave the default Disable radio button selected to prevent initial fragments from being used. This option is not valid for rules that match L4 information such as TCP port number, because that information is carried in the initial packet.

- **Flow Label**: The **Flow Label** field is enabled only if the selection from the **Protocol Type** menu is ICMPv6. The flow label is 20-bit number that is unique to an IPv6 packet and that is used by end stations to signify quality-of-service handling in routers. The range for the flow label is from 0 to 1048575.
- **IPv6 DSCP Service**: Specify the IP DiffServ Code Point (DSCP) field. This is an optional configuration. The DSCP is defined as the high-order six bits of the service type octet in the IPv6 header. Enter an integer from 0 to 63. To select the IPv6 DSCP, select one of the DSCP keywords. To specify a numeric value, select Other and enter the numeric value of the DSCP.
- **Rate Limit Conform Data Rate**: Specify the value of the conforming data rate, which is the data rate at which confirming traffic is limited. The range is from 1 to 4294967295 Kbps.
- **Rate Limit Burst Size**: Specify the value of the burst size, which is the size at which bursts of traffic above the confirming data rate are permitted. The range is from 1 to 128 Kbps.
- **Time Range**: From the **Time Range** menu, select the timer schedule that must be associated with the rule. For more information about timer schedules, see Timer schedules on page 153. If you did not set up any timer schedules, the menu does not present any options.

6. Click the **Apply** button.
   Your settings are saved.
   The Rule Status field in the table shows if the rule is active or inactive. Blank means that no timer schedule is associated with the rule.
7. To save the settings to the running configuration, click the **Save** icon.

Modify the match criteria for an IPv6 ACL rule

You can modify the match criteria for an IPv6 ACL rule.

To modify the match criteria for an IPv6 ACL rule:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IPv6 Rules**.
   The IPv6 Rules page displays.
3. From the **ACL Name** menu, select the ACL that includes the rule that you want to modify.
4. In the IPv6 ACL Rule Table, click the rule.
   The rule is a hyperlink. The IPv6 ACL Rule Configuration page displays.
5. Modify the IPv6 ACL rule criteria.

6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.

Delete an IPv6 ACL rule

You can delete an IPv6 ACL rule that you no longer need.

To delete an IPv6 ACL rule:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IPv6 Rules**.
   The IPv6 Rules page displays.
3. From the **ACL Name** menu, select the ACL that includes the rule that you want to delete.
4. In the IPv6 ACL Rule Table, select the check box that is associated with the rule.
5. Click the **Delete** button.
   The rule is removed.
6. To save the settings to the running configuration, click the **Save** icon.

Configure IP ACL interface bindings

When you bind a basic IPv4, extended IPv4, or IPv6 ACL to an interface, all the rules that you defined for the IP ACL are applied to the selected interface.
If resources on the switch are insufficient, an attempt to bind an ACL to an interface fails.

To bind an IP ACL to one or more interfaces:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > IP Binding Configuration**.
   The IP Binding Configuration page displays.
3. From the **ACL ID** menu, select an IP ACL.

NOTE: Binding an ACL to an interface fails is the switch has no resources to bind a new ACL. In addition, you cannot bind both an IPv4 ACL and an IPv6 ACL to an interface.

4. From the **Direction** menu, select the packet filtering direction for the ACL:
   - **Inbound**: The rule is applied to traffic entering the port.
   - **Outbound**: The rule is applied to traffic leaving the port.
5. In the **Sequence Number** field, optionally specify a number to indicate the order of the access list relative to other access lists already assigned to the interface and direction.
   A low number indicates high precedence order. If a sequence number is already in use for the interface and direction, the specified access list replaces the currently attached access list using that sequence number. If you do not specify the sequence number (meaning that the value is 0), a sequence number that is one number greater than the highest sequence number currently in use for the interface and direction is used. The range is from 1 to 4294967295.
6. To add the selected ACL to a port or LAG, in the Ports table (or if a stack is configured, in one of the Ports tables) or LAG table, click the port or LAG so that a check mark displays.
   The Ports and LAG tables display the available interfaces for ACL bindings. You can add the ACL to several ports and LAGs.
   The ports for the switch (Unit 1) are displayed. If a stack is configured, the ports for each stacked switch (Unit 1, Unit 2, and so on) are displayed. In addition, all LAGs are displayed, whether or not a stack is configured.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 194.Interface Binding Status information

Field	Description
Interface	The selected interface.
Direction	The selected packet filtering direction for the ACL.
ACL Type	The type of ACL assigned to the selected interface and direction.
ACL ID/Name	The ACL number or name identifying the ACL assigned to the selected interface and direction.
Sequence Number	The sequence number signifying the order of specified ACL relative to other ACLs assigned to the selected interface and direction.

Display or delete IP ACL bindings in the IP ACL binding table

You can display or delete bindings for basic IPv4, extended IPv4, and IPv6 ACLs.

To display or delete IP ACL bindings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > Binding Table**.
   The IP ACL Binding Table page displays. The table displays the IP ACL-to-interface bindings.
3. To delete an IP ACL-to-interface binding, do the following:
   a. Select the check box next to the interface.
   b. Click the **Delete** button.
      The binding is removed.
4. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 195.Interface Binding Status information

Field	Description
Interface	The selected interface.
Direction	The selected packet filtering direction for the ACL.
ACL Type	The type of ACL assigned to the selected interface and direction.
ACL ID/Name	The ACL number or name identifying the ACL assigned to the selected interface and direction.
Sequence Number	The sequence number signifying the order of specified ACL relative to other ACLs assigned to the selected interface and direction.

Configure VLAN ACL bindings

You can associate a MAC ACL, any type of IPv4 ACL, or an IPv6 ACL with a VLAN. When you do so, the ACL is applied to all interfaces that are members of the VLAN.

Add a VLAN ACL binding

You can bind a VLAN ID to an ACL and apply the packet filtering direction for the ACL to either inbound traffic or outbound traffic.

To add a VLAN ACL binding:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL> Advanced > VLAN Binding Table**.
   The VLAN Binding Table page displays.
3. In the **VLAN ID** field, enter the VLAN ID to which the binding must apply.
4. From the **Direction** menu, select the packet filtering direction for the ACL:
   - **Inbound**: The rule is applied to traffic entering the port.
   - **Outbound**: The rule is applied to traffic leaving the port.
5. In the **Sequence Number** field, optionally specify a number to indicate the order of the access list relative to other access lists already assigned to the VLAN and direction.
   A low number indicates high precedence order. If a sequence number is already in use for the VLAN ID and direction, the specified access list replaces the currently attached access list using that sequence number. If you do not specify the sequence number (meaning that the value is 0), a sequence number that is one number greater than the highest sequence number currently in use for the VLAN ID and direction is used. The range is from 1 to 4294967295.
6. From the **ACL Type** menu, select the type of ACL.
   You can select a MAC ACL, IP ACL, or IPv6 ACL.
7. From the **ACL ID** list, select the ID or name of the ACL that must be bound to the specified VLAN.
8. Click the **Add** button.
   The VLAN ACL binding is added.
9. To save the settings to the running configuration, click the **Save** icon.

Remove a VLAN ACL binding

You can remove a binding between a VLAN ID and an ACL.

To remove a VLAN ACL binding:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > ACL > Advanced > VLAN Binding Table**.
   The VLAN Binding Table page displays.
3. Select the check box for the VLAN binding that you want to remove.

4. Click the **Delete** button.
   The VLAN ACL binding is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Monitor the Switch and Network

This chapter covers the following topics:

- Port and EAP packet statistics
- Perform a cable test
- Logs
- Syslog and log server host settings
- Trap log
- Event log
- Configure USB logging
- Port mirroring
- RSPAN VLANs and source and destination switches
- sFlow monitoring

Port and EAP packet statistics

You can view port statistics, including detailed statistics, and Extensible Authentication Protocol (EAP) packets statistics.

Display or clear port statistics

You can display a summary of per-port traffic statistics on the switch and clear the statistics.

To view or clear port statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Ports > Port Statistics**.
   The Port Statistics page displays.
3. Select whether to display physical interfaces, LAGs, or both by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **All**: Both physical interfaces and LAGs are displayed, or for a switch stack, both physical interfaces on all switches in the stack and LAGs are displayed.

4. To find a single interface, select the check box associated with the port, or type the port number in the **Go To Interface** field and click the **Go** button.
5. To clear counters, do one of the following:
   - To clear all the counters for all ports on the switch, select the check box in the row heading and click the **Clear** button.
   - To clear the counters for a specific port, select the check box for the port and click the **Clear** button.
6. To refresh the page, click the **Refresh** button.
7. To save the settings to the running configuration, click the **Save** icon.

The following table describes the per-port statistics displayed on the page.

Table 196.Port statistics information

Field	Description
Interface	The interface.
Total Packets Received Without Errors	The total number of packets that were received without errors.
Packets Received With Error	The number of inbound packets that contained errors, preventing them from being delivered to a higher-layer protocol.
Broadcast Packets Received	The number of good packets received that were directed to the broadcast address. This does not include multicast packets.
Packets Transmitted Without Errors	The number of frames that were transmitted by this port to its segment.
Transmit Packet Errors	The number of outbound packets that could not be transmitted because of errors.
Collision Frames	The best estimate of the total number of collisions on this Ethernet segment.
Link Down Events	The total number of link down events on the port.
Link Flaps	The number of occurrences of link-down-to-link up events (that is, one link flap) during the debouncing time.
Received Rate (Mbps)	The received data rate in Mbps.
Transmitted Rate (Mbps)	The transmitted data rate in Mbps.
Received Error Rate	The received data rate with errors in Mbps.
Transmitted Error Rate	The transmitted data rate with errors in Mbps.
Packets Received Per Second	The number of received packets per second.
Packets Transmitted Per Second	The number of transmitted packets per second.
Time since counters last cleared	The elapsed time in days, hours, minutes, and seconds since the statistics for this port were last cleared.

Display or clear detailed statistics for a port

For a specific port, you can display a variety of per-port traffic statistics or clear the statistics.

To display and clear detailed statistics for a port:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Ports > Port Detailed Statistics**.
   The Port Detailed Statistics page displays.
3. From the **Interface** menu, select the interface (a physical port or LAG) for which you want to display the statistics.
4. From the **MST ID** menu, select the MST ID associated with the interface (if available).
5. To refresh the page, click the **Refresh** button.
6. To clear all the counters, click the **Clear** button.
   All statistics for the port are reset to the default values.
7. To save the settings to the running configuration, click the **Save** icon.

The following table describes the detailed port statistics displayed on the page.

Table 197.Port detailed statistics information

Field	Description
ifIndex	The ifIndex of the interface table entry associated with the port.
Port Type	The port is either Normal or one of the following: • **Mirrored**: The port is a mirrored port in a port mirroring configuration. • **Probe**: The port is the probe port in a port mirroring configuration. • **Port Channel**: The port is a member of a LAG.
Port Channel ID	If the port is a member of a port channel, the port channel's interface ID and name display. Otherwise, Disable displays.
Port Role	Each MST bridge port that is enabled is assigned a port role for each spanning tree. The port role is one of the following values: Root, Designated, Alternate, Backup, Master, or Disabled.
STP Mode	The Spanning Tree Protocol administrative mode that is associated with the port or port channel. The options Enable (spanning tree is enabled (spanning tree is enabled for the port) or Disable (spanning tree is disabled for the port).
STP State	The port's current spanning tree state. This state controls what action a port takes on receipt of a frame. If the bridge detects a malfunctioning port it places that port into the broken state. The states are defined in IEEE 802.1D: • Disabled • Blocking • Listening • Learning • Forwarding • Broken
Admin Mode	The port control administration state. The port must be enabled for it to be allowed into the network. The default is enabled.
Flow Control Mode	Indicates if flow control is enabled or disabled for the port. This field does not apply to LAGs.
LACP Mode	The Link Aggregation Control Protocol administrative state. The mode must be enabled for the port to participate in link aggregation.
Physical Mode	The port speed and duplex mode. In autonegotiation mode the duplex

Field	Description
	mode and speed are set by the autonegotiation process.
Physical Status	The port speed and duplex mode.
Link Status	Indicates if the link is up or down.
Link Trap	Indicates if the port sends a trap when link status changes.
Packets RX and TX 64 Octets	The total number of packets (including bad packets) received or transmitted that were 64 octets in length (excluding framing bits but including FCS octets).
Packets RX and TX 65-127 Octets	The total number of packets (including bad packets) received or transmitted that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 128-255 Octets	The total number of packets (including bad packets) received or transmitted that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 256-511 Octets	The total number of packets (including bad packets) received or transmitted that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 512-1023 Octets	The total number of packets (including bad packets) received or transmitted that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 1024-1518 Octets	The total number of packets (including bad packets) received or transmitted that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 1519-2047 Octets	The total number of packets (including bad packets) received or transmitted that were between 1519 and 2047 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 2048-4095 Octets	The total number of packets (including bad packets) received or transmitted that were between 2048 and 4095 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 4096-9216 Octets	The total number of packets (including bad packets) received or transmitted that were between 4096 and 9216 octets in length inclusive (excluding framing bits but including FCS octets).
Octets Received	The total number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets). This object can be used as a reasonable estimate of Ethernet

Field	Description
	utilization. If greater precision is desired, the etherStatsPkts and etherStatsOctets objects must be sampled before and after a common interval.
Packets Received 64 Octets	The total number of packets (including bad packets) received that were 64 octets in length (excluding framing bits but including FCS octets).
Packets Received 65-127 Octets	The total number of packets (including bad packets) received that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Received 128-255 Octets	The total number of packets (including bad packets) received that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Received 256-511 Octets	The total number of packets (including bad packets) received that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Received 512-1023 Octets	The total number of packets (including bad packets) received that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Received 1024-1518 Octets	The total number of packets (including bad packets) received that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Received > 1518 Octets	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.
Total Packets Received Without Errors	The total number of packets received that were without errors.
Unicast Packets Received	The number of subnetwork-unicast packets delivered to a higher-Layer protocol.
Multicast Packets Received	The total number of good packets received that were directed to a multicast address. This number does not include packets directed to the broadcast address.
Broadcast Packets Received	The total number of good packets received that were directed to the broadcast address. This does not include multicast packets.
Receive Packets Discarded	The number of inbound packets that were discarded even though no errors were detected that would prevent the packets from being

Field	Description
	delivered to a higher-layer protocol. A possible reason for discarding a packet could be to free up buffer space.
Total Packets Received with MAC Errors	The total number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
Jabbers Received	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad frame check sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (alignment error). This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.
Fragments Received	The total number of packets received that were less than 64 octets in length with ERROR CRC (excluding framing bits but including FCS octets).
Undersize Received	The total number of packets received that were less than 64 octets in length with GOOD CRC (excluding framing bits but including FCS octets).
Alignment Errors	The total number of packets received with a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad frame check sequence (FCS) with a nonintegral number of octets.
Rx FCS Errors	The total number of packets received with a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad frame check sequence (FCS) with an integral number of octets
Overruns	The total number of frames discarded because this port was overloaded with incoming packets, and could not keep up with the inflow.
Total Received Packets Not Forwarded	A count of valid frames received that were discarded (that is, filtered) by the forwarding process.
802.3x Pause Frames Received	A count of MAC control frames received on this interface with an opcode indicating the PAUSE operation. This counter does not increment when the interface is operating in half-duplex mode.

Field	Description
Unacceptable Frame Type	The number of frames discarded from this port due to being an unacceptable frame type.
Total Packets Transmitted (Octets)	The total number of octets of data (including those in bad packets) transmitted on the network (excluding framing bits but including FCS octets). This object can be used as a reasonable estimate of Ethernet utilization. If greater precision is desired, the etherStatsPkts and etherStatsOctets objects must be sampled before and after a common interval.
Packets Transmitted 64 Octets	The total number of packets (including bad packets) received that were 64 octets in length (excluding framing bits but including FCS octets).
Packets Transmitted 65-127 Octets	The total number of packets (including bad packets) received that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Transmitted 128-255 Octets	The total number of packets (including bad packets) received that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Transmitted 256-511 Octets	The total number of packets (including bad packets) received that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Transmitted 512-1023 Octets	The total number of packets (including bad packets) received that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Transmitted 1024-1518 Octets	The total number of packets (including bad packets) received that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Transmitted > 1518 Octets	The total number of packets transmitted that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed. This counter has a max increment rate of 815 counts per sec at 10 Mb/s.
Maximum Frame Size	The maximum Ethernet frame size that the interface supports or is configured to use, including the Ethernet header, CRC, and payload. The minimum size is 1500; The maximum size is 12270; The default size is 9198.
Total Packets Transmitted	The number of frames that were transmitted to the port's segment.
Unicast Packets	The total number of packets that higher-level protocols requested be

Field	Description
Transmitted	transmitted to a subnetwork-unicast address, including those that were discarded or not sent.
Multicast Packets Transmitted	The total number of packets that higher-level protocols requested be transmitted to a multicast address, including those that were discarded or not sent.
Broadcast Packets Transmitted	The total number of packets that higher-level protocols requested be transmitted to the broadcast address, including those that were discarded or not sent.
Transmit Packets Discarded	The total number of outbound packets that were discarded even though no errors were detected that would prevent the packets from being delivered to a higher-layer protocol. A possible reason for discarding a packet could be to free up buffer space.
Total Transmit Errors	The sum of single, multiple, and excessive collisions.
Tx FCS Errors	The total number of transmitted packets with a length (excluding framing bits but including FCS octets) between 64 and 1518 octets, inclusive, and with a bad Frame Check Sequence (FCS) with an integral number of octets.
Total Transmit Packets Discarded	The sum of single collision frames discarded, multiple collision frames discarded, and excessive frames discarded.
Single Collision Frames	The number of successfully transmitted frames on a particular interface for which transmission is inhibited by exactly one collision.
Multiple Collision Frames	The number of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collision.
Excessive Collision Frames	The number of frames for which transmission on a particular interface fails due to excessive collisions.
STP BPDUs Received	The number of STP BPDUs received.
STP BPDUs Transmitted	The number of STP BPDUs transmitted.
RSTP BPDUs Received	The number of RSTP BPDUs received.
RSTP BPDUs Transmitted	The number of RSTP BPDUs transmitted.
MSTP BPDUs Received	The number of MSTP BPDUs received.
MSTP BPDUs Transmitted	The number of MSTP BPDUs transmitted.

Field	Description
802.3x Pause Frames Transmitted	The number of MAC control frames transmitted with an opcode indicating the PAUSE operation. This counter does not increment when the interface is operating in half-duplex mode.
GVRP PDUs Received	The number of GVRP PDUs received in the GARP layer.
GVRP PDUs Transmitted	The number of GVRP PDUs transmitted from the GARP layer.
GVRP Failed Registrations	The number of times attempted GVRP registrations could not be completed.
GMRP PDUs Received	The number of GMRP PDUs received from the GARP layer.
GMRP PDUs Transmitted	The number of GMRP PDUs transmitted from the GARP layer.
GMRP Failed Registrations	The number of times attempted GMRP registrations could not be completed.
EAPOL Frames Received	The number of valid EAPOL frames of any type that were received by this authenticator.
EAPOL Frames Transmitted	The number of EAPOL frames of any type that were transmitted by this authenticator.
Load Interval	The period in seconds for which data is used to compute load statistics.
Received Rate (Mbps)	The received data rate in Mbps.
Transmitted Rate (Mbps)	The transmitted data rate in Mbps.
Received Error Rate	The received data rate with errors in Mbps.
Transmitted Error Rate	The transmitted data rate with errors in Mbps.
Packets Received Per Second	The number of received packets per second.
Packets Transmitted Per Second	The number of transmitted packets per second.
Percent Utilization Received	The utilization for received traffic, rounded to the nearest whole percentage.
Percent Utilization Transmitted	The utilization for transmitted traffic, rounded to the nearest whole percentage.
Time Since Counters Last Cleared	The elapsed time in days, hours, minutes, and seconds since the statistics for this port were last cleared.

Display or clear EAP and EAPoL statistics

You can display information about Extensible Authentication Protocol (EAP) and EAP over LAN (EAPoL) packets that are received on physical ports.

To display or clear EAP and EAPoL statistics:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Ports > EAP Statistics**.
   The EAP Statistics page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.
4. To clear the counters, which resets the EAP and EAPoL statistics to default values, take one of the following actions:
   - To clear the counters for a specific port, select the check box associated with the port, and click the **Clear** button.
   - To clear the counters for multiple ports, select the check boxes associated with the ports, and click the **Clear** button.
   - To clear all counters for all ports, select the check box in the row heading, and click the **Clear** button.
5. To save the settings to the running configuration, click the **Save** icon.
6. To refresh the page, click the **Refresh** button.

The following table describes the EAP statistics displayed on the page.

Table 198.EAP Statistics information

Field	Description
Port	The port number.
PAE Capabilities	The PAE capabilities of the port.
EAPOL Frames Received	The number of valid EAPOL frames of any type that were received by this authenticator.
EAPOL Frames Transmitted	The number of EAPOL frames of any type that were transmitted by this authenticator.
EAPOL Start Frames Received	The number of EAPOL start frames that were received by this authenticator.
EAPOL Logoff Frames Received	The number of EAPOL logoff frames that were received by this authenticator.
EAPOL Last Frame Version	The protocol version number carried in the most recently received EAPOL frame.
EAPOL Last Frame Source	The source MAC address carried in the most recently received EAPOL frame.
EAPOL Invalid Frames Received	The number of EAPOL frames that were received by this authenticator in which the frame type is not recognized.
EAPOL Length Error Frames Received	The number of EAPOL frames that were received by this authenticator in which the frame type is not recognized.
EAP Response/ID Frames Received	The number of EAP response/identity frames that were received by this authenticator.
EAP Response Frames Received	The number of valid EAP response frames (other than resp/ID frames) that were received by this authenticator.
EAP Request/ID Frames Transmitted	The number of EAP request/identity frames that were transmitted by this authenticator.
EAP Request Frames Transmitted	The number of EAP request frames (other than request/identity frames) that were transmitted by this authenticator.

Perform a cable test

You can test and display information about the cables that are connected to switch ports.

To perform a cable test:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Ports > Cable Test**.
   The Cable Test page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.
4. Select the check boxes that are associated with the physical ports for which you want to test the cables.
5. Click the **Apply** button.
   A cable test is performed on the selected interface. The cable test might take up to two seconds to complete. If the port has an active link, the cable status is always Normal. The command returns a cable length estimate if this feature is supported by the PHY for the current link speed. Note that if the link is down and a cable is attached to a 10/100 Ethernet adapter then the cable status might be Open or Short because some Ethernet adapters leave unused wire pairs unterminated or grounded.

The following table describes the view-only fields on the page.

Table 199.Cable Test

Field	Description
Cable Status	Indicates the cable status: • **Normal**: The cable is working correctly. • **Open**: The cable is disconnected or a faulty connector exists. • **Short**: An electrical short exists in the cable. • **Cable Test Failed**: The cable status could not be determined. The cable might in fact be working. • **Untested**: The cable is not yet tested. • **Invalid cable type**: The cable type is unsupported.
Cable Length	The estimated length of the cable in meters. The length is displayed as a range between the shortest estimated length and the longest estimated length. Unknown is displayed if the cable length could not be determined. The cable length is displayed only if the cable status is Normal.
Failure Location	The estimated distance in meters from the end of the cable to the failure location. The failure location is displayed only if the cable status is Open or Short.

Logs

The switch generates messages in response to events, faults, and errors as well as changes in the configuration or other occurrences. These messages are stored locally and can be forwarded to one or more centralized points of collection for monitoring purposes or long-term archival storage (see Syslog and log server host settings on page 743). Local and remote configuration of the logging capability includes filtering of messages logged or forwarded based on severity and generating component.

Manage and display the memory log

The memory log stores messages in memory based upon the settings for message component and severity. You can set the administrative status and behavior of logs in the system buffer. These log messages are cleared when the switch reboots.

To manage and display the memory log:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Monitoring > Logs > Memory Log**.
   The Memory Log page displays.
3. Select one of the following Admin Status radio buttons:
   - **Enable**: Enable the switch to log default messages. This is the default setting.
   - **Disable**: Prevent the switch from logging default messages.
4. From the **Behavior** menu, specify the behavior of the log when it is full:
   - **Wrap**: When the buffer is full, the oldest log messages are deleted as the system logs new messages.
   - **Stop on Full**: When the buffer is full, the system stops logging new messages and preserves all existing log messages.
5. From the **Severity Filter** menu, select the logging level for messages that must be logged. Log messages with the selected severity level and all log messages of greater severity are logged. For example, if you select **Warning**, the logged messages include Warning, Error, Critical, Alert, and Emergency. The default severity level is Informational (6).
   The severity can be one of the following levels:
   - **Emergency**: Level 0, the highest warning level. If the device is down or not functioning properly, an emergency log is saved to the device.
   - **Alert**: Level 1, the second-highest warning level. An alert log is saved if a serious device malfunction occurs, such as all device features being down.
   - **Critical**: Level 2, the third-highest warning level. A critical log is saved if a critical device malfunction occurs, for example, two device ports are not functioning, while the rest of the device ports remain functional.
   - **Error**: Level 3, a device error occurred, such as a port being offline.
   - **Warning**: Level 4, the lowest level of a device warning.
   - **Notice**: Level 5, provides the network administrators with device information.
   - **Informational**: Level 6, provides device information. This is the default setting.
   - **Debug**: Level 7, provides detailed information about the log.

   NOTE: A log records messages equal to or above a configured severity threshold.
6. In the **Threshold** field, enter the percentage of log space that, if exceeded, causes logging to stop.
   Enter a percentage from 1 to 100. The default is 80.
   The threshold applies only if the selection from the **Behavior** menu is **Stop-on-Full**. By default, the default the selection from the **Behavior** menu is **Wrap** and the threshold does not apply.

7. Click the **Apply** button.
   Your settings are saved.
   The Memory Log table displays on the Memory Log page.
   The Total number of Messages field displays the number of messages the system logged in memory. By default, up to 160 (80 percent of the maximum of 200) of the most recent entries can be displayed on the page.
   The rest of the page displays the Memory Log messages. The format of the log message is the same for messages that are displayed for the message log, persistent log, or console log. Messages logged to a collector or relay through syslog support the same format as well.
   The following example shows the standard format for a log message:
   ```
   <189> Jun 21 2024 00:00:18: AAA-5-CONNECT: New http connection for user admin, source 192.168.1.111 ACCEPTED
   ```
   The message was generated as severity 189 (Notice, level 5) on June 21, 2024 at 00:00:18 a.m. by component AAA. The message indicates that the administrator successfully logged on to the HTTP management interface from a host with IP address 192.168.1.111.
   The severity number in the message is generated by a combination of the facility and the severity. In the previous example, the default facility of 23 is multiplied by 8 and the severity (5) is added: (23 * 8) + 5 = 189
8. To refresh the page, click the **Refresh** button.
9. To clear the log, click the **Clear** button.
10. To save the settings to the running configuration, click the **Save** icon.

Message log format

This topic applies to the format of all logged messages that are displayed for the message log, persistent log, or console log.
Messages logged to a collector or relay through syslog use an identical format:

- <15>Aug 24 05:34:05 0.0.0.0-1 MSTP[2110]: mspt_api.c(318) 237%% Interface 12 transitioned to root state on message age timer expiry. This example indicates a message with severity 7 (15 mod 8) (debug) on a switch and generated by component MSTP running in thread ID 2110 on Aug 24 05:34:05 by line 318 of file mstp_api.c. This is the 237th message logged with system IP 0.0.0.0 and task-ID 1.
- <15>Aug 24 05:34:05 STK0 MSTP[2110]: mspt_api.c(318) 237%% Interface 12 transitioned to root state on message age timer expiry. This example indicates a user-level message (1) with severity 7 (debug) on a system that is not a switch and generated by component MSTP running in thread ID 2110 on Aug 24 05:34:05 by line 318 of file mstp_api.c. This is the 237th message logged. Messages logged to a collector or relay through syslog use a format identical to the previous message.
- Total number of Messages: For the message log, only the latest 200 entries are displayed on

the page.

Enable or disable the command configuration log

The switch can log changes that are made with CLI configuration commands.

To enable or disable the command configuration log:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Logs > Command Log Configuration**.
   The Command Log Configuration page displays.
3. Select one of the following Admin Status radio buttons:
   - **Enable**: Enable the switch to log CLI configuration commands. This is the default setting.
   - **Disable**: Prevent the switch from logging CLI configuration commands.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.

Enable or disable console logging

The switch can log messages to a serial device that is connected to the switch.

To enable or disable console logging:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Logs > Console Log Configuration**.
   The Console Log Configuration page displays.
3. Select one of the following Admin Status radio buttons:
   - **Enable**: Enable the switch to log messages to a serial device. This is the default setting.
   - **Disable**: Prevent the switch from logging messages to a serial device.

4. From the **Severity Filter** menu, select the logging level for messages that must be logged. Log messages with the selected severity level and all log messages of greater severity are logged. For example, if you select **Warning**, the logged messages include Warning, Error, Critical, Alert, and Emergency. The default severity level is Error (3).
   The severity can be one of the following levels:
   - **Emergency**: Level 0, the highest warning level. If the device is down or not functioning properly, an emergency log is saved to the device.
   - **Alert**: Level 1, the second-highest warning level. An alert log is saved if a serious device malfunction occurs, such as all device features being down.
   - **Critical**: Level 2, the third-highest warning level. A critical log is saved if a critical device malfunction occurs, for example, two device ports are not functioning, while the rest of the device ports remain functional.
   - **Error**: Level 3, a device error occurred, such as a port being offline. This is the default setting.
   - **Warning**: Level 4, the lowest level of a device warning.
   - **Notice**: Level 5, provides the network administrators with device information.
   - **Informational**: Level 6, provides device information.
   - **Debug**: Level 7, provides detailed information about the log.

   NOTE: A log records messages equal to or above a configured severity threshold.

5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Syslog and log server host settings

You can let the switch send log messages to one or more servers, that is, to remote logging hosts. A remote log server is the same as a remote syslog host.
You must enable the server log on the switch and specify one or more remote syslog host.

Configure the syslog settings

To configure the syslog settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Logs > Server Log**.
   The Server Log page displays.

3. Select one of the following Admin Status radio buttons:
   - **Enable**: The switch sends log messages to all configured syslog servers.
   - **Disable**: The switch stops sending log messages all configured syslog servers.
4. In the **Local UDP Port** field, specify the port number that the switch uses to send log messages to the syslog servers.
   The range is from 1 to 65535. The default port is 514.
5. From the **Source Interface** menu, select the source interface or VLAN that must be used for syslog communication. By default, the following options display in the menu:
   - **None**: The primary IP address of the originating (outbound) interface is used as the source address.
   - **VLAN 1**: The primary IP address of VLAN 1 is used as the source address. This is the default selection.
   - **Service Port**: The management port IP address is used as the source address.

   Depending on the configuration of your switch, the following options can display:
   - **Another VLAN ID**: The primary IP address of a VLAN other than VLAN 1 is used as the source address.
   - **Routing interface**: The primary IP address of a routing interface is used as the source address.
   - **Routing VLAN**: The primary IP address of a VLAN routing interface is used as the source address.
   - **Routing loopback interface**: The primary IP address of a routing loopback interface is used as the source address.
   - **Different**: For some features, *Different* can display. This means that the source interface is configured separately.

   NOTE: The syslog function on the switch can save the log messages simultaneously to a remote server and to a USB storage device.
6. To simultaneously save the log messages to a USB storage device, in the **USB Filename** field, specify the name of the USB file.
   The file name cannot include the following symbols: V:*?”<>!. Up to 64 characters can be entered. The 64 characters are only the file name length, the extension is automatically added.
7. Click the **Apply** button.
   Your settings are saved.
8. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 200.Syslog configuration information

Field	Description
Messages Received	The number of messages received by the log process. This includes messages that are dropped or ignored.
Messages Relayed	The number of syslog messages relayed.
Messages Ignored	The number of syslog messages ignored.

Add a syslog server

To add a syslog server:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Logs > Server Log**.
   The Server Log page displays.

3. In the Server Configuration table, configure the following settings:
   - **IP Address Type**: Specify the IP address type of the syslog server, which can be **IPv4**, **IPv6**, or **DNS**.
   - **Host Address**: Specify the IP address or host name of the syslog server.
   - **Port**: Specify the port on the syslog server. The default port number is 514.
   - **Severity Filter**: Use the menu to select the severity of the logs that must be sent to the syslog server. Logs with the selected severity level and all logs of greater severity are sent to the host. For example, if you select **Error**, the logged messages include Error, Critical, Alert, and Emergency. The severity can be one of the following levels:
     - **Emergency**: Level 0, the highest warning level. If the device is down or not functioning properly, an emergency log is saved to the device.
     - **Alert**: Level 1, the second-highest warning level. An alert log is saved if a serious device malfunction occurs, such as all device features being down.
     - **Critical**: Level 2, the third-highest warning level. A critical log is saved if a critical device malfunction occurs, for example, two device ports are not functioning, while the rest of the device ports remain functional.
     - **Error**: Level 3, a device error occurred, such as a port being offline.
     - **Warning**: Level 4, the lowest level of a device warning.
     - **Notice**: Level 5, provides the network administrators with device information.
     - **Informational**: Level 6, provides device information.
     - **Debug**: Level 7, provides detailed information about the log.
4. Click the **Add** button.
   The syslog server is added.
   The Status field in the Server Configuration table shows whether the syslog server is currently active.
5. To save the settings to the running configuration, click the **Save** icon.

Modify the settings for a syslog server

To modify the settings for a remote syslog host:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Logs > Server Log**.
   The Server Log page displays.
3. Select the check box that is associated with the syslog server.
4. Change the settings as needed.

5. Click the **Apply** button.
   Your settings are saved.

Delete the settings for a syslog server

To delete the settings for a remote syslog host:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Logs > Server Log**.
   The Server Log page displays.
3. Select the check box that is associated with the syslog server.
4. Click the **Delete** button.
   The syslog server is removed.

Trap log

The trap log includes information about the traps that the switch sent. You can display and clear the entries in the trap log.
You can also retrieve the trap log and save it as a file. For more information, see Export a file from the switch on page 767.

To display or clear the trap log:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Logs > Trap Logs**.
   The Trap Logs page displays.
3. To refresh the page, click the **Refresh** button.
4. To clear the log, click the **Clear** button.
5. To save the settings to the running configuration, click the **Save** icon.

The following table describes the view-only fields on the page.

Table 201. Trap log information

Field	Description
Number of Traps Since Last Reset	The number of traps that occurred since the switch last rebooted.
Trap Log Capacity	The maximum number of traps stored in the log. If the number of traps exceeds the capacity, the new entries overwrite the oldest entries.
Number of Traps since log last viewed	The number of traps that occurred since the traps were last displayed. **NOTE:** If you save the trap log to another device (that is, you upload the file from the switch), this counter is set to 0.
Log	The sequence number of this trap.
System Up Time	The time when this trap occurred, expressed in days, hours, minutes and seconds, since the last reboot of the switch.
Trap	Information about the trap.

Event log

You can display and clear the event log, which contains error messages from the switch.

NOTE: If you reset the switch to factory default settings, the event log is not cleared.

To display or clear the event log:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Logs > Event Logs**.
   The Event Logs page displays.
3. To refresh the page, click the **Refresh** button
4. To clear the log, click the **Clear** button.
5. To save the settings to the running configuration, click the **Save** icon.

The following table describes the event log information displayed on the page.

Table 202.Event logs information

Field	Description
Entry	The sequence number of the event.
Type	The type of the event.
File Name	The file in which the event originated.
Line	The line number of the event.
Task ID	The task ID of the event.
Code	The event code.
Time	The time this event occurred.

Configure USB logging

If you connect a USB storage device such as a flash drive to the switch, the switch logs messages to the device.

To configure USB logging:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Logs > USB Log Configuration**.
   The USB Log Configuration page displays.
3. From the **USB Log Mode** menu, select how logging occurs:
   - **Overwrite on full**: The switch logs messages to the USB device and overwrites old messages if the USB is full. This is the default setting.
   - **Stop on Full**: The switch logs messages to the USB device and stops logging if the USB device is full.
   - **Disable**: Prevent the switch from logging messages to the USB device.

4. From the **USB Log Mode** menu, select the logging level for messages that must be logged. Log messages with the selected severity level and all log messages of greater severity are logged. For example, if you select **Warning**, the logged messages include Warning, Error, Critical, Alert, and Emergency. The default severity level is Error (3).
   The severity can be one of the following levels:
   - **Emergency**: Level 0, the highest warning level. If the device is down or not functioning properly, an emergency log is saved to the device.
   - **Alert**: Level 1, the second-highest warning level. An alert log is saved if a serious device malfunction occurs, such as all device features being down.
   - **Critical**: Level 2, the third-highest warning level. A critical log is saved if a critical device malfunction occurs, for example, two device ports are not functioning, while the rest of the device ports remain functional.
   - **Error**: Level 3, a device error occurred, such as a port being offline.
   - **Warning**: Level 4, the lowest level of a device warning.
   - **Notice**: Level 5, provides the network administrators with device information. This is the default setting.
   - **Informational**: Level 6, provides device information.
   - **Debug**: Level 7, provides detailed information about the log.

   NOTE: A log records messages equal to or above a configured severity threshold.

5. Click the **Apply** button.
   Your settings are saved.
   The USB Log Operational Status field displays is logging is enabled (Active or Inactive).
6. To save the settings to the running configuration, click the **Save** icon.

Port mirroring

Port mirroring lets you select the network traffic of specific switch ports for analysis by a network analyzer. For each port mirroring configuration (the switch supports up to four), you can select many switch ports as source ports but only a single switch port as the destination port. For each port mirroring configuration, you can configure how traffic is mirrored on a source port by selecting packets that are received, transmitted, or both. You can also apply an IP ACL or MAC ACL to filter the traffic that is being mirrored.
A packet that is copied to the destination port is in the same format as the original packet. This means that if the mirror is copying a received packet, the copied packet is VLAN-tagged or untagged as it was received on the source port. If the mirror is copying a transmitted packet, the copied packet is VLAN-tagged or untagged as it is being transmitted on the source port.

To set up a port mirroring configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Mirroring > Multiple Port Mirroring**.
   The Multiple Port Mirroring page displays.
3. From the **Session ID** menu, select **1**, **2**, **3**, or **4** as the port mirroring session ID.
   You can set up four port mirroring configurations that functions independently from each other.
4. Enable or disable the port mirroring configuration by selecting one of the following Admin Mode radio buttons:
   - **True**: The selected source interfaces are mirrored. That is, traffic entering or leaving the source interfaces is copied (mirrored) to the selected destination port.
   - **False**: The selected source interfaces are not mirrored.
5. From the **Destination Port** menu, select the destination interface to which traffic must be copied.
   For each port mirroring configuration, you can configure one destination port only. The default is None.
6. From the **Filter Type** menu, select if traffic must be filtered:
   - **None**: All traffic is mirrored.
   - **IP ACL**: Traffic is filtered according to the IP ACL for which you must type the name in the **Filter Name** field. Only traffic that matches the rules in the IP ACL is mirrored to the destination port.
     For more information about IP ACLs, see Configure a basic or extended IPv4 ACL on page 699.
   - **MAC ACL**: Traffic is filtered according to the MAC ACL for which you must type the name in the **Filter Name** field. Only traffic that matches the rules in the MAC ACL is mirrored to the destination port.
     For more information about MAC ACLs, see Configure a MAC ACL on page 691.

   The following steps refer to the Source Interface Configuration section, in which you must select one or more source interfaces and set the traffic direction.

7. Select whether to display physical interfaces, LAGs, the CPU, VLANs, or all by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **CPU**: Only the CPU is displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Physical interfaces, LAGs, the CPU, and VLANs are displayed, or for a switch stack, physical interfaces on all switches in the stack, LAGs, the CPU, and VLANs are displayed.
8. Use one of the following methods to select an interface:
   - In the **Go To Interface** field, enter the interface in the unit/slot/port format and click on the **Go** button.
   - Next to the Interface column, select the check box for the interface or interfaces that you want to include.

   NOTE: If an interface is a member of a VLAN *and* is member of a LAG, you cannot select the VLAN as a source VLAN. However, if an interface is a member of a VLAN that you select as a source VLAN, the interface *can* be a member of a LAG.
9. From the **Direction** menu, select the direction of the traffic that must be mirrored from the configured source ports.
   - **None**: The port (or VLAN) is not mirrored.
   - **Tx and Rx**: Both egress (transmitted) and egress (received) traffic is mirrored. For a VLAN, this option select the VLAN as the source VLAN.
   - **Rx**: Ingress (received) traffic only is mirrored. This option is not available for a VLAN.
   - **Tx**: Egress (transmitted) traffic only is mirrored. This option is not available for a VLAN.
10. Click the **Apply** button.
    Your settings are saved.
    If the port is configured as a source port, the Status field displays Mirrored.
    If the port is configured as a destination port, the Status field displays Probe.
11. To save the settings to the running configuration, click the **Save** icon.

RSPAN VLANs and source and destination switches

You can configure a VLAN as a remote switched port analyzer (RSPAN) VLAN. RSPAN lets you mirror traffic from multiple source ports (or from all ports that are members of a VLAN) from

different network devices and send the mirrored traffic to a destination port (a probe port connected to a network analyzer) on a remote device. The mirrored traffic is tagged with the RSPAN VLAN ID and transmitted over trunk ports in the RSPAN VLAN.
A switch can function as an RSPAN source switch that sends traffic to another switch that functions as an RSPAN destination switch. Because the switch supports multiple RSPAN configurations, the switch can also function as an RSPAN destination switch that receives traffic from another switch that functions as an RSPAN source switch.

Configure an existing VLAN as an RSPAN VLAN

You can configure an existing VLAN as an RSPAN VLAN.

To configure an existing VLAN as an RSPAN VLAN:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Mirroring > RSPAN VLAN**.
   The RSPAN VLAN page displays.
3. Select the check box for the VLAN that must become an RSPAN VLAN.
4. From the **Admin Mode** menu, select **Enable** to configure the VLAN as an RSPAN VLAN.
   By default, the selection is Disable and the VLAN is not an RSPAN VLAN.
5. Click the **Apply** button.
   Your settings are saved.
6. To save the settings to the running configuration, click the **Save** icon.

Configure the switch as an RSPAN source switch

You can configure the switch as an RSPAN source switch that sends traffic to another switch that functions as an RSPAN destination switch.
You can create up to four RSPAN source switch configurations.

To configure the switch as an RSPAN source switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Mirroring > RSPAN Source Switch Configuration**.
   The RSPAN Source Switch Configuration page displays.

3. From the **Session ID** menu, select **1**, **2**, **3**, or **4** as the RSPAN source switch session ID.
   You can set up four RSPAN source switch configurations that functions independently from each other.
4. Enable or disable the RSPAN source switch configuration by selecting one of the following Admin Mode radio buttons:
   - **True**: The selected source interfaces are mirrored. That is, traffic entering or leaving the source interfaces is copied (mirrored) to the selected RSPAN destination VLAN.
   - **False**: The selected source interfaces are not mirrored.
5. From the **RSPAN Destination VLAN** menu, select the destination VLAN to which traffic must be copied.
   For each RSPAN source switch configuration, you can configure one RSPAN destination VLAN port only. The default is None.
6. From the **RSPAN Reflector Port** menu, select the interface that is connected to the RSPAN destination switch.
7. From the **Filter Type** menu, select if traffic must be filtered:
   - **None**: All traffic is sent to the destination VLAN.
   - **IP ACL**: Traffic is filtered according to the IP ACL for which you must type the name in the **Filter Name** field. Only traffic that matches the rules in the IP ACL is sent to the destination VLAN.
     For more information about IP ACLs, see Configure a basic or extended IPv4 ACL on page 699.
   - **MAC ACL**: Traffic is filtered according to the MAC ACL for which you must type the name in the **Filter Name** field. Only traffic that matches the rules in the MAC ACL is sent to the destination VLAN.
     For more information about MAC ACLs, see Configure a MAC ACL on page 691.

   The following steps refer to the RSPAN Source Interface Configuration section, in which you must select one or more source interfaces and set the traffic direction.

8. Select whether to display physical interfaces, LAGs, the CPU, VLANs, or all by clicking one of the following links above the table heading:
   - **1 or Unit ID for a stacked switch**:
     - **1**: If no switch stack is configured, the physical interfaces for the switch are displayed.
     - **Unit ID for a stacked switch**: If a switch stack is configured, the physical interfaces for the switch with the selected stack unit ID are displayed.
   - **LAG**: Only LAGs are displayed.
   - **CPU**: Only the CPU is displayed.
   - **VLANS**: Only VLANs are displayed.
   - **All**: Physical interfaces, LAGs, the CPU, and VLANs are displayed, or for a switch stack, physical interfaces on all switches in the stack, LAGs, the CPU, and VLANs are displayed.
9. Use one of the following methods to select an interface:
   - In the **Go To Interface** field, enter the interface in the unit/slot/port format and click on the **Go** button.
   - Next to the Interface column, select the check box for the interface or interfaces that you want to include.

   NOTE: If an interface is a member of a VLAN *and* is member of a LAG, you cannot select the VLAN as a source VLAN. However, if an interface is a member of a VLAN that you select as a source VLAN, the interface *can* be a member of a LAG.
10. From the **Direction** menu, select the direction of the traffic that must be mirrored from the configured source ports.
    - **None**: The port (or VLAN) is not mirrored.
    - **Tx and Rx**: Both egress (transmitted) and egress (received) traffic is mirrored. For a VLAN, this option select the VLAN as the source VLAN.
    - **Rx**: Ingress (received) traffic only is mirrored. This option is not available for a VLAN.
    - **Tx**: Egress (transmitted) traffic only is mirrored. This option is not available for a VLAN.
11. Click the **Apply** button.
    Your settings are saved.
    If the port is configured as a source port, the Status field displays Mirrored.
12. To save the settings to the running configuration, click the **Save** icon.

Configure the switch as an RSPAN destination switch

You can configure the switch as an RSPAN destination switch that receives traffic from another switch that functions as an RSPAN source switch.
You can create up to four RSPAN destination switch configurations.

To configure the switch as an RSPAN destination switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring > Mirroring > RSPAN Destination Switch Configuration**.
   The RSPAN Destination Switch Configuration page displays.
3. From the **Session ID** menu, select **1**, **2**, **3**, or **4** as the RSPAN destination switch session ID.
   You can set up four RSPAN destination switch configurations that functions independently from each other.
4. Enable or disable the RSPAN destination switch configuration by selecting one of the following Admin Mode radio buttons:
   - **True**: Traffic from the selected RSPAN source VLAN is copied (mirrored) to the selected RSPAN destination port.
   - **False**: Traffic from the selected RSPAN source VLAN is not mirrored.
5. From the **RSPAN Source VLAN** menu, select the RSPAN source VLAN from which traffic must be copied.
6. From the **RSPAN Destination Port** menu, select the RSPAN destination port to which traffic must be copied.
   For each RSPAN destination switch configuration, you can configure one RSPAN destination port only. The default is None.
7. From the **Filter Type** menu, select if traffic must be filtered:
   - **None**: All traffic is sent to the RSPAN destination port.
   - **IP ACL**: Traffic is filtered according to the IP ACL for which you must type the name in the **Filter Name** field. Only traffic that matches the rules in the IP ACL is sent to the RSPAN destination port.
     For more information about IP ACLs, see Configure a basic or extended IPv4 ACL on page 699.
   - **MAC ACL**: Traffic is filtered according to the MAC ACL for which you must type the name in the **Filter Name** field. Only traffic that matches the rules in the MAC ACL is sent to the RSPAN destination port.
     For more information about MAC ACLs, see Configure a MAC ACL on page 691.

8. Click the **Apply** button.
   Your settings are saved.
9. To save the settings to the running configuration, click the **Save** icon.

sFlow monitoring

sFlow is a standard for monitoring high-speed switched and routed networks. sFlow technology is built into the switch and many other network devices to give visibility into network activity.
The sFlow monitoring system consists of an embedded sFlow agent and a central sFlow collector, also referred to as an sFlow receiver. The sFlow agent uses sampling technology to capture traffic statistics from the device it is monitoring. sFlow datagrams are used to forward the sampled traffic statistics to an sFlow collector for analysis.
The sFlow agent uses two forms of sampling: statistical packet-based sampling of switched or routed packet flows, and time-based sampling of counters.

sFlow agent overview

Packet flow sampling and counter sampling are performed by sFlow instances associated with individual data sources within the sFlow agent. Packet flow samples and counter samples are combined in sFlow datagrams. Packet flow sampling causes a steady but random stream of sFlow datagrams to be sent to the sFlow collector. Counter samples can be taken opportunistically to fill these datagrams.
To perform packet flow sampling, you must configure an sFlow sampler instance with a sampling Rate. To perform counter sampling, you must configure an sFlow poller instance with a polling interval. The sFlow agent sends the collected information in the form of sFlow datagrams to sFlow collectors to the sFlow receiver.

Configure the source interface for the sFlow agent

The switch can function as an sFlow agent and forward polled and sampled traffic to an sFlow receiver. You can configure the source interface that connects to the device that functions as the sFlow receiver.

To configure the source interface for the sFlow agent:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring> sFlow > Basic > sFlow Agent Information**.
   The sFlow Agent Information page displays.

3. From the **Source Interface** menu, select the management interface that the sFlow agent must use. By default, the following options display in the menu:
   - **None**: The primary IP address of the originating (outbound) interface is used as the source address.
   - **VLAN 1**: The primary IP address of VLAN 1 is used as the source address. This is the default selection.
   - **Service Port**: The management port IP address is used as the source address.

   Depending on the configuration of your switch, the following options can display:
   - **Another VLAN ID**: The primary IP address of a VLAN other than VLAN 1 is used as the source address.
   - **Routing interface**: The primary IP address of a routing interface is used as the source address.
   - **Routing VLAN**: The primary IP address of a VLAN routing interface is used as the source address.
   - **Routing loopback interface**: The primary IP address of a routing loopback interface is used as the source address.
   - **Different**: For some features, *Different* can display. This means that the source interface is configured separately.
4. Click the **Apply** button.
   Your settings are saved.
5. To save the settings to the running configuration, click the **Save** icon.
   The following table describes the view-only fields on the page.

Table 203.sFlow agent information

Field	Description
Agent Version	The version and implementation of the MIB. The agent version consists of the following components: • MIB version: For example, 1.3. • Organization: NETGEAR • Revision: For example, 13.0.2.10.
Agent Address	The IP address associated with the agent. The IP address depends on your selection of the source interface.

Configure an sFlow receiver

You can configure one or more devices that must function as sFlow receivers. The sFlow source interface provides the connection to these devices.

To configure an sFlow receiver:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring> sFlow > Advanced > sFlow Receiver Configuration**.
   The sFlow Receiver Configuration page displays.
3. Select the **Receiver Index** check box for a receiver.
4. In the **Receiver Owner** field, specify the receiver owner.
   Specify the name of the owner or the device.
   By default, the time-out option is enabled but the time-out period is configured as 0, which means that sFlow sampling does not stop after it is enabled.
5. To configure a time-out period, do the following:
   a. In the **Receiver Timeout** field, specify the time-out period in seconds after which sampling stops.
      The range is from 0 to 2147483647. A value of zero means that sFlow sampling does not stop after it is enabled.
   b. From the **No Timeout** menu, select **False** to enable the time-out sampling option.
      The default is False. If you select **True**, sampling does not time out and the time-out period is automatically set to the maximum of 2147483647.
6. In the **Maximum Datagram Size** field, specify the maximum number of data bytes that can be sent in a single sample datagram.
   The default value is 1400. The range is from 200 to 12188. To prevent fragmentation of the datagrams, we recommend that you do not set the datagram size too low.
7. In the **Receiver Address** field, specify the IP address of the sFlow receiver.
8. In the **Receiver Port** field, specify the destination port for sFlow datagrams.
   The default port number is 6343. The range is from 1 to 65535.
9. Click the **Apply** button.
   Your settings are saved.
   The Receiver Datagram Version field displays 5 as the version of sFlow datagrams.
10. To save the settings to the running configuration, click the **Save** icon.

Configure sFlow polling and sampling on an interface

If the switch is configured as an sFlow agent, the switch can do the following:

- Collect statistical packet-based sampling of switched flows and sends them to the configured receivers. A data source that is configured to collect flow samples is called a sampler.
- Collect time-based sampling of network interface statistics and send them to the configured sFlow receivers. A data source configured to collect counter samples is called a poller.

You can configure sampling and a polling settings on a switch interface. Traffic on the interface is then sampled and polled and forwarded to the sFlow receiver.

To configure sFlow polling and sampling on an interface:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Monitoring> sFlow > Advanced > sFlow Interface Configuration**.
   The sFlow Interface Configuration page displays.
3. If a stack is configured, select whether to display the physical interfaces for one switch or for all switches in the stack:
   - **Unit ID for a stacked switch**: The physical interfaces for the switch with the selected stack unit ID are displayed. If no switch stack is configured, the only option is unit ID 1.
   - **All**: The physical interfaces for all switches in the stack are displayed. If no switch stack is configured, the All option does not have any effect.
4. Use one of the following methods to select an interface for the flow poller and sampler:
   - In the **Go To Interface** field, enter the interface in the unit/slot/port format and click on the **Go** button.
   - Next to the Interface column, select the check box for the interface that you want to use.

   You can select physical ports only.
5. In the **Poller Receiver Index** field, specify the sFlow receiver that must be associated with the poller.
   For information about sFlow receivers, see Configure an sFlow receiver on page 758.
6. In the **Poller Interval** field, specify the number of seconds between successive polling.
   A sampling interval of 0 disables sampling. The range is from 0 to 86400 seconds. The default is 0 seconds.
7. In the **Sampler Receiver Index** field, specify the sFlow receiver that must be associated with the sampler.
   For information about sFlow receivers, see Configure an sFlow receiver on page 758.

8. In the **Sampling Rate** field, specify the statistical sampling rate for packet sampling.
   A sampling rate of 1 counts all packets. A sampling rate of 0 disables sampling. The allowed range is 1024 to 65536. The default is 0.
9. In the **Maximum Header Size** field, specify the maximum number of bytes to be copied from a sampled packet.
   The range is from 20 to 256 bytes. The default is 128 bytes.
10. Click the **Apply** button.
    Your settings are saved.
11. To save the settings to the running configuration, click the **Save** icon.

Maintenance and Troubleshooting

This chapter covers the following topics:

- Save the configuration
- Automatic installation of the configuration file
- Reboot the switch from the main UI
- Reset the switch to the factory default settings
- Export a file from the switch
- Update software or download a file
- Manage software images
- Diagnostics and troubleshooting

Save the configuration

When you save the configuration, changes that you made are retained by the switch when it is rebooted. You can also save the settings on each configuration page of the main UI by clicking the **Save** icon.

To save the configuration:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Save Config > Save Configuration**.
   The Save Configuration page displays.
3. Select the check box.
4. Click the **Apply** button.
   Your settings are saved. If you restart the switch, the saved settings are retained.

Automatic installation of the configuration file

If the switch is initialized without a configuration file present on the switch, the automatic installation (auto install) process can automatically load the switch configuration from a server in the network.

Configure the auto install process

The auto install process requires that DHCP is enabled by default on the switch. To save the downloaded configuration to the startup configuration, you must explicitly save the file as such.
The auto Install process depends on the configuration of other devices in the network, including a DHCP or BOOTP server, a TFTP server and, if necessary, a DNS server.
The auto install process occurs in three phases:

1. **Configuration or assignment of an IP address to the switch**: For more information, see IPv4 management interfaces and VLANs on page 49 or IPv6 management interfaces and VLANs on page 54.
2. **Assignment of a TFTP server**: If the switch functions as a DHCP client (which it does by default), the DHCP server can automatically provide the TFTP server to the switch, for example, if the DHCP server is configured with option 125 (see Option 125 DHCP server requirements for obtaining an configuration file through auto install on page 765).

3. **Downloading a configuration file for the switch from the TFTP server**: See the following procedure.

   To configure the auto install process:

   a. Launch a web browser and enter the IP address of the switch in the address field of your web browser.
   The login page displays.

   > **NOTE:** The default web browser protocol is HTTPS. Do not use HTTP to access the UI.

   b. Click the **Main UI Login** button.
   The main UI login page displays in a new tab.
   c. Enter **admin** as the user name and enter one of the following passwords:
      - Enter your local device password. The first time that you log in, no password is required. However, you then must specify a local device password to use each subsequent time that you log in.
      - If you are managing the switch through the Engage controller, enter the site password for the Engage site to which the switch is onboarded.
      - If, in addition to the main UI, you are also managing the switch through the Insight Cloud Portal or Insight app, enter the Insight network password for the Insight network to which the switch is added.
   d. Click the **Login** button.
   The System Information page displays.
   e. Select **Maintenance > Save Config > Auto Install Configuration**.
   The Auto Install Configuration page displays.
   f. From the **AutoInstall Mode** menu, select to start or stop the auto install process:
      - **Start**: Starts the auto install process (after you click the **Apply** button).
      - **Stop**: Stops the auto install process (after you click the **Apply** button).
   g. From the **AutoInstall Persistent Mode** menu, select to enable or disable the persistent mode for the auto install process:
      - **Enabled**: The auto install process saves the downloaded configuration file and applies it after the switch reboots. This is the default setting.
      - **Disabled**: The auto install process does not save the downloaded configuration file and, therefore, cannot apply it after the switch reboots.
   h. From the **AutoSave Mode** menu, select to enable or disable automatic saving of the configuration to the startup configuration:
      - **Enabled**: The automatically downloaded configuration is saved to the startup configuration.
      - **Disabled**: The automatically downloaded configuration is not saved to the startup configuration. If you want to save the configuration you must do to manually. This is

the default setting.

i. In the **AutoInstall Retry Count** menu, specify the number of times that the switch sends a unicast message to the TFTP server.
The TFTP server is specified in the DHCP configuration (see Option 125 DHCP server requirements for obtaining an configuration file through auto install on page 765). After the specified number of times is exceeded, the switch sends broadcast messages to the TFTP server. The range is from 1 to 3. By default, the number is 3.

j. Click the **Apply** button.
Your settings are saved.
The Auto-Install State field displays the status of the auto install process.

Option 125 DHCP server requirements for obtaining an configuration file through auto install

You can use auto install to let the switch download a configuration file (image file) from a TFTP server using DHCP option 125. The image update can either upgrade or downgrade the firmware on the switch.
For the switch to be able to download an image from a TFTP server, the DHCP server must include an image description file that lists the image that the switch must download from the TFTP server. For example, the `autoinstall_dhcp` image description file on the DHCP server lists the `M4XXX-v1.2.3.4.stk` image that the switch must download from the TFTP server.
Option 125 in the DHCP server must contain the following information (the example uses the `autoinstall_dhcp` image description file and the `M4XXX-v1.2.3.4.stk` image):

- **Enterprise number (4 octets)**. 0x0000 0x11ae In decimals, the enterprise number for NETGEAR is 4526.
- **Data length (1 octet)**: 0x12 The data length includes the SubOption code plus the SubOption length plus the image description file (not the image itself): 1+1+16.
  - **SubOption code (2 octets)**: 0x05
  - **SubOption length (1 octet)**: 0x10 The length of the name of the image description file.
  - **Image description file name (N octets)**: 61.7574.6f69.6e73.7461.6c6c.5f64.6863.70 In plain text, the name of the image description file is `autoinstall_dhcp`.

The content of the `autoinstall_dhcp`image description file is `M4XXX-v1.2.3.4.stk`, which is the image that the switch must download from the TFTP server.
You can change the name for the image description file, but then you must also change the option 125 data length, which is based on the SubOption code, SubOption length, and image description file name.

Reboot the switch from the main UI

You can reboot the switch from the main UI. If your configuration includes a stack of switches, you can select which switch to reboot, or you can reboot all switches in the stack.

NOTE: If you can physically access the switch, you can reboot the switch by pressing the multi-function **Reset** button on the front panel for less than 5 seconds. (Do not press the button for more than 5 seconds!)

To reboot the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Reset > Device Reboot**.
   The Device Reboot page displays.
3. From the Reboot Unit No. menu, select which switch to reboot, or select **All** to reboot all switches in a stack.
   If your switch does not function in a stack, select **1**.
4. Select one of the following radio buttons:
   - **Save prior to reboot**: The switch saves all changes that you made and for which you did not click the **Save** icon in the main interface.
   - **Don't save prior to reboot**: The switch does not save the changes that you made and for which you did not click the **Save** icon in the main interface.
5. Click the **Apply** button.
   An Alert pop-up window displays.
6. Click the **OK** button to confirm.
   The switch reboots.

Reset the switch to the factory default settings

You can reset the switch configuration to the factory default values. All changes that you made are erased. If the IP address changes, your web session might disconnect.

NOTE: If you reset the switch to the default configuration, the management IP address is reset to 169.254.100.100, and the DHCP client is enabled. No IP address is set for the OOB port, but its DHCP client is enabled. If you lose network connectivity after you reset the switch to the factory defaults, see Log in to the main UI with a web browser on page 3.

To reset the switch to the factory default settings:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Reset > Factory Default**.
   The Factory Default page displays.
3. Select the check box.
4. Click the **Apply** button.
   A confirmation pop-up window displays.
5. Click the **OK** button to confirm.
   The configuration is reset to the factory default settings.

Export a file from the switch

You can export configuration (ASCII) or log (ASCII log) files from the switch to a file server over a TFTP, SFTP, SCP, or FTP session, to a computer over an HTTP session, or to a USB storage device.

Export a file from the switch to a server

You can upload (export) configuration (ASCII or log ASCII) and other types of files from the switch to a TFTP, SFTP, SCP, or FTP server on the network.

To export a file from the switch to a server:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.

2. Select **Maintenance > Export > File Export**.
   The File Export page displays.
3. From the **File Type** menu, select the type of file that must be exported:
   - **Text Configuration**: A text-based configuration file enables you to edit a configured text file (`startup-config`) offline as needed. The most common usage of text-based configuration is to upload a working configuration from a device, edit it offline to personalize it for another similar device (for example, change the device name or IP address), and download it to that device. This is the default setting.
   - **Error Log**: The switch error persistent log, also referred to as the event log.
   - **Buffered Log**: The system buffered (in-memory) log.
   - **Trap Log**: The trap log with the switch trap records.
   - **Script File**: The script file that you must specify in the **Local File Name** field.
   - **CLI Banner**: The CLI banner file.
   - **Tech Support**: The tech support file is a text-based file that contains a variety of hardware, software, and configuration information that can assist in device and network troubleshooting.
   - **Crash Logs**: The switch crash logs, if any are available.
   - **Backup Configuration**: The backup configuration file is a text-based file.
   - **CPU Packets Capture File**: The CPU packets capture file with the captured and stored CPU packets.
   - **Factory Default Configuration**: The factory default configuration file is a text-based file with the stored factory default configuration.
   - **DHCP Client Configuration File**: The DHCP client configuration file with the DHCP server (host) information.
4. From the **Transfer Mode** menu, specify the protocol that must be used to transfer the file to the server:
   - **TFTP**: Trivial File Transfer Protocol.
   - **SFTP**: Secure File Transfer Protocol. This selection requires you to specify a user name and password.
   - **SCP**: Secure Copy Protocol. This selection requires you to specify a user name and password.
   - **FTP**: File Transfer Protocol. This selection requires you to specify a user name and password.

5. From the **Server Address Type** menu, select the format for the **Server Address** field:
   - **IPv4**: Indicates that the server address is an IPv4 address in dotted-decimal format. This is the default setting.
   - **IPv6**: Indicates that the server address is an IPv6 address in hexadecimal numbers using 16-bit values between colons.
   - **DNS**: Indicates that the server address is a host name.
6. In the **Server Address** field, enter the IP address of the server in accordance with the format indicated by the server address type.
   The default is the IPv4 address 0.0.0.0.
7. In the **Remote File Path** field, specify the path on the server where you want to save the file. The path name can include alphabetic, numeric, forward slash, dot, or underscore characters only. You can enter up to 160 characters.
8. In the **Remote File Name** field, specify a destination file name for the file to be exported. You can enter up to 32 characters.
9. If you select **Script File** from the **File Type** menu, specify the name of the script file in the **Local File Name** field.
10. If you select **SFTP**, **SCP**, or **FTP** from **Transfer Mode** menu, specify the following settings:
    - **User Name**: Specify the user name for remote login to the server.
    - **Password**: Specify the password for remote login to the server.
11. Click the **Apply** button.
    The file is exported (uploaded) to the server. The page displays information about the progress of the file transfer.

Use HTTP to export a file from the switch to a computer

You can upload (export) files of various types from the switch to a computer through an HTTP session on your web browser.

To use HTTP to export a file from the switch to a computer:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Export > HTTP File Export**.
   The HTTP File Export page displays.

3. From the **File Type** menu, select the type of file that must be exported:
   - **Text Configuration**: A text-based configuration file enables you to edit a configured text file (`startup-config`) offline as needed. The most common usage of text-based configuration is to upload a working configuration from a device, edit it offline to personalize it for another similar device (for example, change the device name or IP address), and download it to that device. This is the default setting.
   - **Error Log**: The switch error persistent log, also referred to as the event log.
   - **Buffered Log**: The system buffered (in-memory) log.
   - **Trap Log**: The trap log with the switch trap records.
   - **Script File**: The script file that you must specify in the **Local File Name** field.
   - **CLI Banner**: The CLI banner file.
   - **Tech Support**: The tech support file is a text-based file that contains a variety of hardware, software, and configuration information that can assist in device and network troubleshooting.
   - **Crash Logs**: The switch crash logs, if any are available.
   - **Backup Configuration**: The backup configuration file is a text-based file.
   - **CPU Packets Capture File**: The CPU packets capture file with the captured and stored CPU packets.
   - **Factory Default Configuration**: The factory default configuration file is a text-based file with the stored factory default configuration.
   - **DHCP Client Configuration File**: The DHCP client configuration file with the DHCP server (host) information.
4. If you select **Script File** from the **File Type** menu, specify the name of the script file in the **Local File Name** field.
5. Click the **Apply** button.
   The file is exported (uploaded) to the computer. The page displays information about the progress of the file transfer.

Export a file from the switch to a USB storage device

You can upload (export) a text configuration file or DHCP client configuration file to a USB storage device that is attached to the switch.

To export a file from the switch to a USB storage device:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Export > USB File Export**.
   The USB File Export page displays.

3. From the **File Type** menu, select the type of file that must be exported:
   - **Text Configuration**: A text-based configuration file enables you to edit a configured text file (`startup-config`) offline as needed. The most common usage of text-based configuration is to upload a working configuration from a device, edit it offline to personalize it for another similar device (for example, change the device name or IP address), and download it to that device. This is the default setting.
   - **DHCP Client Configuration File**: The DHCP client configuration file with the DHCP server (host) information.
4. In the **File Path** field, specify the path on the USB storage device.
   You can use up to 146 characters. The default is blank.
5. In the **USB File** field, specify a name for the file. This is the name under which the file will be saved on the USB storage device.
   You can enter up to 32 characters.
6. Click the **Apply** button.
   The file is exported (uploaded) to the USB storage device. The page displays information about the progress of the file transfer.

Update software or download a file

You can download firmware, configuration (ASCII), log (ASCII log), SSL, PEM, or other types of files to the switch from a file server over a TFTP, SFTP, SCP, or FTP session, from a computer over an HTTP session, or from a USB storage device.
You can manually check for the latest software version (also referred to as firmware version), download the firmware to a server or computer, and then download the firmware to the switch. If firmware release notes are available with new firmware, read the release notes to find out if you must reconfigure the switch after updating.
In this context, downloading is also referred to as upgrading.
Note the following about SSH and SSL files:

- **SSH**: For you to be able to download SSH files to the switch, SSH must be administratively disabled and no active SSH sessions must occur.
- **SSL**: SSL files contain information to encrypt, authenticate, and validate HTTPS sessions. For you to be able to download SSL files to the switch, HTTPS must be administratively disabled.

Download a software file or another type of file from a server to the switch

You can download a software (firmware) image and configuration, SSL, PEM, or other types of files from an TFTP, SFTP, SCP, or FTP server on your network to the switch.
Before you download a file to the switch, the following conditions must be true:

- The file to download from the server is on the server in the appropriate directory.
- The file is in the correct format.
- The switch contains a path to the server.

To download a file from a server to the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Upgrade > File Upgrade**.
   The File Upgrade page displays.

3. From the **File Type** menu, select the type of file:
   - **Software**: The system software image, which is saved in one of two flash sectors called images (image1 and image2). The active image stores the active copy, while the other image stores a second copy. The device boots and runs from the active image. If the active image is corrupted, the system automatically boots from the nonactive image. This is a safety feature for faults occurring during the boot upgrade process. This is the default setting.
   - **Text Configuration**: A text-based configuration file enables you to edit a configured text file (`startup-config`) offline as needed. The most common usage of text-based configuration is to upload a working configuration from a device, edit it offline to personalize it for another similar device (for example, change the device name or IP address), and download it to that device.
   - **SSH-2 RSA Key PEM File**: An SSH-2 Rivest-Shamir-Adelman (RSA) key file (PEM Encoded).
   - **SSH-2 DSA Key PEM File**: An SSH-2 Digital Signature Algorithm (DSA) key file (PEM Encoded).
   - **ECDSA Key PEM File**: An Elliptic Curve Digital Signature Algorithm (ECDSA) key File (PEM Encoded).
   - **SSL Trusted Root Certificate PEM File**: An SSL Trusted Root Certificate file (PEM Encoded).
   - **SSL Server Certificate PEM File**: An SSL Server Certificate file (PEM Encoded) to the device.
   - **SSL DH 1024 Encryption Parameter PEM File**: An SSL Diffie-Hellman 1024-bit Encryption Parameter File (PEM Encoded) to the device.
   - **SSL DH 2048 Encryption Parameter PEM File**:An SSL Diffie-Hellman 2048-bit Encryption Parameter File (PEM Encoded) to the device.
   - **Script File**: A text-based configuration script file. You must use the command-line interface (CLI) to configure, validate, and activate the script.
   - **CLI Banner**: The CLI banner contains the text to be displayed on the CLI before the login prompt.
   - **IAS Users**: The Internal Authentication Server (IAS) users database file. The IAS user database stores a list of user name and (optional) password values for local port-based user authentication.
   - **Factory Default Configuration**: The factory default configuration file is a text-based file with the stored factory default configuration.
   - **Public Key Configuration**: The public key file used for configuration script validation.
   - **Application**: Files for an application.
   - **Tech Support Commands File**: A file with tech support commands.
   - **DHCP Client Configuration File**: The DHCP client configuration file with the DHCP server (host) information.

4. If the selection from the **File Type** menu is **Software**, the Image Name menu displays and you must select the software image that must be downloaded to the switch:
   - **Image1**: Select image1 to upload image1.
   - **Image2**: Select image2 to upload image2.

   We recommended that you do not overwrite the active image.
5. If the selection from the **File Type** menu is **Software**, the Verify radio buttons display so that you can select one of the following options:
   - **None**: Verification of the downloaded file is disabled. This is the default setting.
   - **Verify**: The digital signature of the downloaded file is verified.
   - **No Verify**: The digital signature of the downloaded file is not verified.
6. If the selection from the **File Type** menu is **Application**, the **Application File Name** field displays and you must specify the application file name.
7. From the **Transfer Mode** menu, specify the protocol that must be used to transfer the file to the server:
   - **TFTP**: Trivial File Transfer Protocol.
   - **SFTP**: Secure File Transfer Protocol. This selection requires you to specify a user name and password.
   - **SCP**: Secure Copy Protocol. This selection requires you to specify a user name and password.
   - **FTP**: File Transfer Protocol. This selection requires you to specify a user name and password.
8. From the **Server Address Type** menu, select the format for the **Server Address** field:
   - **IPv4**: Indicates that the server address is an IPv6 address in dotted-decimal format. This is the default setting.
   - **IPv6**: Indicates that the server address is an IPv6 address in hexadecimal numbers using 16-bit values between colons.
   - **DNS**: Indicates that the server address is a host name.
9. In the **Server Address** field, enter the IP address of the server in accordance with the format indicated by the server address type.
   The default is the IPv4 address 0.0.0.0.
10. In the **Remote File Path** field, specify the path on the server where the file is located.
    The path name can include alphabetic, numeric, forward slash, dot, or underscore characters only. You can enter up to 160 characters.
11. In the **Remote File Name** field, specify the file name for the file to be downloaded.
    You can enter up to 32 characters.

12. If you select **SFTP**, **SCP**, or **FTP** from **Transfer Mode** menu, specify the following settings:
    - **User Name**: Specify the user name for remote login to the server.
    - **Password**: Specify the password for remote login to the server.
13. Click the **Apply** button.
    The file is downloaded from the server to the switch. The page displays information about the progress of the file transfer.
14. To save the settings to the running configuration, click the **Save** icon.

Use HTTP to download a software file or another type of file to the switch

You can download files of various types from a computer to the switch through an HTTP session on your web browser.

To use HTTP to download a file from a computer to the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Upgrade > HTTP File Upgrade**.
   The HTTP File Upgrade page displays.

3. From the **File Type** menu, select the type of file:
   - **Software**: The system software image, which is saved in one of two flash sectors called images (image1 and image2). The active image stores the active copy, while the other image stores a second copy. The device boots and runs from the active image. If the active image is corrupted, the system automatically boots from the nonactive image. This is a safety feature for faults occurring during the boot upgrade process. This is the default setting.
   - **Text Configuration**: A text-based configuration file enables you to edit a configured text file (`startup-config`) offline as needed. The most common usage of text-based configuration is to upload a working configuration from a device, edit it offline to personalize it for another similar device (for example, change the device name or IP address), and download it to that device.
   - **SSH-2 RSA Key PEM File**: An SSH-2 Rivest-Shamir-Adelman (RSA) key file (PEM Encoded).
   - **SSH-2 DSA Key PEM File**: An SSH-2 Digital Signature Algorithm (DSA) key file (PEM Encoded).
   - **ECDSA Key PEM File**: An Elliptic Curve Digital Signature Algorithm (ECDSA) key File (PEM Encoded).
   - **SSL Trusted Root Certificate PEM File**: An SSL Trusted Root Certificate file (PEM Encoded).
   - **SSL Server Certificate PEM File**: An SSL Server Certificate file (PEM Encoded) to the device.
   - **SSL DH 1024 Encryption Parameter PEM File**: An SSL Diffie-Hellman 1048-bit Encryption Parameter File (PEM Encoded) to the device.
   - **SSL DH 2048 Encryption Parameter PEM File**:An SSL Diffie-Hellman 2048-bit Encryption Parameter File (PEM Encoded) to the device.
   - **Config Script**: A text-based configuration script file. You must use the command-line interface (CLI) to configure, validate, and activate the script.
   - **CLI Banner**: The CLI banner contains the text to be displayed on the CLI before the login prompt.
   - **IAS Users**: The Internal Authentication Server (IAS) users database file. The IAS user database stores a list of user name and (optional) password values for local port-based user authentication.
   - **Factory Default Configuration**: The factory default configuration file is a text-based file with the stored factory default configuration.
   - **Public Key Image**: The public key file used for code image validation.
   - **DHCP Client Configuration File**: The DHCP client configuration file with the DHCP server (host) information.

4. If the selection from the **File Type** menu is **Software**, the **Image Name** menu displays and you must select the software image that must be downloaded to the switch:
   - **Image1**: Select image1 to upload image1.
   - **Image2**: Select image2 to upload image2.

   We recommended that you do not overwrite the active image.
5. If the selection from the **File Type** menu is **Software**, the Verify radio buttons display so that you can select one of the following options:
   - **None**: Verification of the downloaded file is disabled. This is the default setting.
   - **Verify**: The digital signature of the downloaded file is verified.
   - **No Verify**: The digital signature of the downloaded file is not verified.
6. Next to Select File, click the **Browse** button, navigate to the file, and select the file to download.
   You can select a file with a file name of up to 80 characters.
   The software file is in `.stk` format.
7. Click the **Apply** button.
   The download begins.
   The Download Status field displays the status during transfer file to the switch.

   NOTE: After a file transfer starts, wait until the page refreshes. When the page refreshes, the Select File option is blanked out, indicating that the file transfer is complete.
8. To save the settings to the running configuration, click the **Save** icon.

Download a software file or another type of file from a USB storage device to the switch

You can download a software file, text configuration file, or DHCP client configuration file from a USB storage device to the switch.

To download a file from a USB storage device to the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Upgrade > USB File Upgrade**.
   The USB File Upgrade page displays.

3. From the **File Type** menu, select the type of file that must be downloaded:
   - **Software**: The system software image, which is saved in one of two flash sectors called images (image1 and image2). The active image stores the active copy, while the other image stores a second copy. The device boots and runs from the active image. If the active image is corrupted, the system automatically boots from the nonactive image. This is a safety feature for faults occurring during the boot upgrade process. This is the default setting.
   - **Text Configuration**: A text-based configuration file enables you to edit a configured text file (`startup-config`) offline as needed. The most common usage of text-based configuration is to upload a working configuration from a device, edit it offline to personalize it for another similar device (for example, change the device name or IP address), and download it to that device. This is the default setting.
   - **DHCP Client Configuration File**: The DHCP client configuration file with the DHCP server (host) information.
4. If the selection from the **File Type** menu is **Software**, the **Image Name** menu displays and you must select the software image that must be downloaded to the switch:
   - **Image1**: Select image1 to upload image1.
   - **Image2**: Select image2 to upload image2.

   We recommended that you do not overwrite the active image.
5. If the selection from the **File Type** menu is **Software**, the Verify radio buttons display so that you can select one of the following options:
   - **None**: Verification of the downloaded file is disabled. This is the default setting.
   - **Verify**: The digital signature of the downloaded file is verified.
   - **No Verify**: The digital signature of the downloaded file is not verified.
6. In the **File Path** field, specify the path on the USB storage device.
   You can use up to 146 characters. The default is blank.
7. In the **USB File** field, specify the file name on the USB storage device.
   You can enter up to 32 characters.
8. Click the **Apply** button.
   The file is downloaded from the USB storage device to the switch. The page displays information about the progress of the file transfer.
9. To save the settings to the running configuration, click the **Save** icon.

Download and install an SSL security certificate file on the switch

If you use HTTPS instead of HTTP to access the main UI, you are not required to obtain an SSL certificate. The security warning that might display in your browser prompts you to confirm that the self-signed certificate of the switch is valid. Once you do so, the browser warning might no longer display when you log in.
However, if you obtain an SSL security certificate from a certificate authority, you can download and install the SSL security certificate through an HTTP session using your web browser.
For an SSL security certificate, you must download two Privacy Enhanced Mail (PEM) files to the switch:

- **SSL Trusted Root Certificate PEM File**: The SSL trusted root certificate PEM file, which must be in the format `xxxxCERTxxxxx.pem`.
- **SSL Server Certificate PEM File**: The SSL server certificate PEM file (the key file), which must be in the format `xxxxKEYxxxxx.pem`.

Before you can download and install an SSL security certificate, you must disable HTTPS on the switch.

To disable HTTPS and use an HTTP session to download and install an SSL security certificate file on the switch:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Security > Access > HTTPS > HTTPS Configuration**.
   The HTTPS Configuration page displays.
3. Select the Admin Mode **Disable** radio button.
4. Click the **Apply** button.
   Your settings are saved. Because you changed the access mode from HTTPS to HTTP, you are logged out of the switch.
5. Wait one minute, refresh your browser, and log back in to the switch.
6. Select **Maintenance > Upgrade > HTTP File Upgrade**.
   The HTTP Firmware/File Update page displays.
7. From the **File Type** menu, select **SSL Trusted Root Certificate PEM File**.
8. Select the Select File **Browse** button and locate the file that you want to download.
   This is the certificate file, which must be in the format `xxxxCERTxxxxx.pem`.

9. Click the **Apply** button.
   The file transfer begins.
   The page displays information about the progress of the file transfer. The page refreshes automatically when the file transfer completes (or if it fails).
10. From the **File Type** menu, select **SSL Server Certificate PEM File**.
    This is the key file, which must be in the format `xxxxKEYxxxxx.pem`.
11. Select the Select File **Browse** button and locate the file that you want to download.
    The file name can contain up to 80 characters.
12. Click the **Apply** button.
    The file transfer begins.
    The page displays information about the progress of the file transfer. The page refreshes automatically when the file transfer completes.
13. To save the settings to the running configuration, click the **Save** icon.

Manage software images

The switch maintains two versions of the switch software in permanent storage. One image is the active image, and the second image is the backup image. The active image is loaded when the switch starts or reboots. This feature reduces switch down time when you are upgrading the switch software.

Copy a software image

You can copy a software image from one location (primary or backup) to another.

To copy a software image:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Image Management > Copy**.
   The Copy page displays.
3. Select the Source Image **Image1** or **Image2** radio button to specify the image to be copied.
4. If the switch is a member of a stack, from the **Stack Member** menu, select the switch to which the image must be copied, or or select **All** to copy the image to all switches.
   If your switch does not function in a stack, select **1**.
5. Select the Destination Image **Image1** or **Image2** radio button to specify the destination image.

6. Click the **Apply** button.
   The image is copied.
7. To save the settings to the running configuration, click the **Save** icon.

Configure dual image settings

The Dual Image feature allows the switch to retain two firmware images in permanent storage. You can select which image must be loaded when the reboots, specify an image description, or delete an image. This feature reduces switch down time when you are upgrading or downgrading the software image.

Change the software image that loads when the switch starts or reboots

You can change the software image that loads when the switch starts. If you configured a switch stack, you can select the image that loads for each switch in the stack.

To change the image that loads during the boot process:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Image Management > Dual Image Configuration**.
   The Dual Image Configuration page displays.
3. Select the check box for the image that is *not* the active image (that is, the image for which the Active Image fields shows False) but that you want the switch to run *after* it reboots.
   If you configured a switch stack, you can select the image that loads for each switch in the stack.
4. From the **Next Active Image** menu, select **True**.
5. As an option, specify a name for the selected image by entering one in the **Image Description** field.
6. Click the **Apply** button.
   Your settings are saved.
7. To save the settings to the running configuration, click the **Save** icon.
8. After activating the image, reboot the switch (see Reboot the switch from the main UI on page 766).
   If you do not reboot the switch, it continues running the image shown in the Current-active field until the next time that the switch reboots.

The following table describes the view-only fields on the page.

Table 204.Dual image configuration information

Field	Description
Unit	The switch stack unit ID. If you did not configure a stack, the unit ID is 1.
Image Name	The name of the image. By default, the image names are image1 and image2.
Active Image	The current active image (True or False).
Version	The firmware version of the image.

Delete a software image

You can delete a software image that is not the active image. If you configured a switch stack, you can delete the image that is not the active image for one or more switches in the stack.

To delete a software image:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Image Management > Dual Image Configuration**.
   The Dual Image Configuration page displays.
3. Select the check box for the image that is *not* the active image (that is, the image for which the Active Image fields shows False).
   If you configured a switch stack, you can select the image for one or more switches in the stack.
   You cannot delete the active image. For the current active image, the Active Image field show True.
4. Click the **Delete** button.
   The image is removed.
5. To save the settings to the running configuration, click the **Save** icon.

Diagnostics and troubleshooting

You can send a ping, trace a route, and perform a memory dump.

Ping an IPv4 address

You can configure the switch to send a ping request to a specified IPv4 address. You can use this option to check whether the switch can communicate with a particular IPv4 device. When you send a ping, the switch sends a specified number of ping requests and the results are displayed. If a reply to the ping is received, the following message displays:

```
PING x.y.z.w (x.y.z.w): size data bytes

size bytes from x.y.z.w: seq=0 ttl=xyz

--- x.y.z.w ping statistics ---

count packets transmitted, count packets received, x% packet loss
```

If a reply to the ping is not received, the following message displays:

```
PING x.y.z.w (x.y.z.w): size data bytes

--- x.y.z.w ping statistics ---

count packets transmitted, 0 packets received, 100% packet loss
```

To ping an IPv4 address:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Troubleshooting > Ping IPv4**.
   The Ping IPv4 page displays.
3. In the **IP Address/Host Name** field, enter the IP address or host name of the device that must be pinged.
   The format is x.x.x.x. The maximum number of characters is 255.
4. In the **Count** field, enter the number of echo requests that must be sent.
   The default value is 3. The range is from 1 to 15.
5. In the **Interval (secs)** field, enter the time between ping packets in seconds.
   The default value is 3 seconds. The range is from 1 to 60.
6. In the **Datagram Size** field, enter the size of the ping packet.
   The default value is 0 bytes. The range is from 0 to 13000.

7. From the **Source** menu, as an option, you can select the IP address or interface that must be used to send echo request packets:
   - **None**: The source address of the ping packet is the address of the default egress interface.
   - **IP Address**: The source IP address that must be used when echo request packets are sent. With this selection, the **IP Address** field displays and you must enter the IPv4 address that must be used as the source.
   - **Interface**: The interface that must be used when echo request packets are sent. The **Interface** menu displays, and you must select the interface.
8. Click the **Apply** button.
   The specified address is pinged. The results are displayed below the configurable data in the Results field.

Ping an IPv6 address

You can configure the switch to send a ping request to a specified IPv6 address. You can use this option to check whether the switch can communicate with a particular IPv6 device. When you send a ping, the switch sends a specified number of ping requests and the results are displayed. If a reply to the ping is received, the following message displays:

```
PING x:y::z:w (x:y::z:w): size data bytes

size bytes from x:y::z:w: seq=0 ttl=xyz

--- x:y::z:w ping statistics ---

count packets transmitted, count packets received, x% packet loss
```

If a reply to the ping is not received, the following message displays:

```
PING x:y::z:w (x:y::z:w): size data bytes

--- x:y::z:w ping statistics ---

count packets transmitted, 0 packets received, 100% packet loss
```

To ping an IPv6 address:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Troubleshooting > Ping IPv6**.
   The Ping IPv6 page displays.

3. From the **Ping** menu, select the type of ping:
   - **Global**: Pings a global IPv6 address.
   - **Link Local**: Pings a link-local IPv6 address over a specified interface. With this selection, the **Interface** menu displays, and you must select the interface.
4. In the **IPv6 Address/Hostname** field, enter the IPv6 address or host name of the station that must be pinged.
   The format is xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx. The maximum number of characters is 255.
5. In the **Count** field, enter the number of echo requests that must be sent.
   The range is from 1 to 15. The default value is 3.
6. In the **Interval (secs)** field, enter the time in seconds between ping packets.
   The range is from 1 to 60. The default value is 3.
7. In the **Datagram Size** field, enter the datagram size.
   The valid range is from 0 to 13000. The default value is 0 bytes.
8. From the **Source** menu, as an option, you can select the IP address or interface that must be used to send echo request packets:
   - **None**: The source address of the ping packet is the address of the default egress interface.
   - **IPv6 Address**: The source IP address that must be used when echo request packets are sent. With this selection, the **IPv6 Address** field displays and you must enter the IPv6 address that must be used as the source.
   - **Interface**: The interface that must be used when echo request packets are sent. The **Interface** menu displays, and you must select the interface.
9. Click the **Apply** button.
   The specified address is pinged. The results are displayed below the configurable data in the Results field.

Send an IPv4 traceroute

You can configure the switch to send a traceroute request to a specified IPv4 address or host name. You can use this to discover the paths that packets take to a remote destination. When you send a traceroute, the switch displays the results below the configurable data.
If a reply to the traceroute is received, the following message displays:

```
traceroute to x.y.z.w (x.y.z.w), maxTTL hops max, size byte packets

initTTL x.y.z.w (x.y.z.w) 0.000 ms * 0.000 ms

initTTL+1 x.y.z.w (x.y.z.w) 0.000 ms * 0.000 ms

initTTL+2 x.y.z.w (x.y.z.w) 0.000 ms * 0.000 ms
```

To send an IPv4 traceroute:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Troubleshooting > Traceroute IPv4**.
   The Traceroute IPv4 page displays.
3. In the **IP Address/Hostname** field, enter the IP address or host name of the device for which the path must be discovered.
4. In the **Probes Per Hop** field, enter the number of probes per hop.
   The default value is 3. The range is from 1 to 10.
5. In the **Max TTL** field, enter the maximum time to live (TTL) for the destination.
   The default value is 30. The range is from 1 to 255.
6. In the **Init TTL** field, enter the initial TTL to be used.
   The default value is 1. The range is from 1 to 255.
7. In the **MaxFail** field, enter the maximum number of failures allowed in the session.
   The default value is 5. The range is from 1 to 255.
8. In the **Interval (secs)** field, enter the time between probes in seconds.
   The default value is 3. The range is from 1 to 60.
9. In the **Port** field, enter the UDP destination port for the probe packets.
   The default value is 33434. The range is from 1 to 65535.
10. In the **Size** field, enter the size of the probe packets.
    The default value is 0. The range is from 0 to 39936.

11. From the **Source** menu, as an option, you can select the IP address or interface that must be used to send probe packets:
    - **None**: The source address of the probe packet is the address of the default egress interface.
    - **IP Address**: The source IP address that must be used when probe request packets are sent. With this selection, the **IP Address** field displays and you must enter the IPv4 address that must be used as the source.
    - **Interface**: The interface that must be used when probe request packets are sent. The **Interface** menu displays, but the only available selection from the menu is **Network**.
12. Click the **Apply** button.
    A traceroute request is sent to the specified IP address or host name. The results are displayed below the configurable data in the Results field.

Send an IPv6 traceroute

You can configure the switch to send a traceroute request to a specified IPv6 address or host name. You can use this to discover the paths that packets take to a remote destination. When you send a traceroute, the switch displays the results below the configurable data.
If a reply to the traceroute is received, the following message displays:

```
traceroute to x:y::z:w (x:y::z:w), maxTTL hops max, size byte packets
```

```
initTTL x:y::z:w (x:y::z:w) 0.000 ms * 0.000 ms
```

```
initTTL+1 x:y::z:w (x:y::z:w) 0.000 ms * 0.000 ms
```

```
initTTL+2 x:y::z:w (x:y::z:w) 0.000 ms * 0.000 ms
```

To send an IPv6 traceroute:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Troubleshooting > Traceroute IPv6**.
   The Traceroute IPv6 page displays.
3. In the **IPv6 Address/Host Name** field, enter the IPv6 address or host name of the device for which the path must be discovered.
4. In the **Probes Per Hop** field, enter the number of probes per hop.
   The default value is 3. The range is from 1 to 10.

5. In the **Max TTL** field, enter the maximum time to live (TTL) for the destination.
   The default value is 30. The range is from 1 to 255.
6. In the **Init TTL** field, enter the initial TTL to be used.
   The default value is 1. The range is from 1 to 255.
7. In the **MaxFail** field, enter the maximum number of failures allowed in the session.
   The default value is 5. The range is from 1 to 255.
8. In the **Interval (secs)** field, enter the time between probes in seconds.
   The default value is 3. The range is from 1 to 60.
9. In the **Port** field, enter the UDP destination port for the probe packets.
   The default value is 33434. The range is from 1 to 65535.
10. In the **Size** field, enter the size of the probe packets.
    The default value is 0. The range is from 0 to 39936.
11. From the **Source** menu, as an option, you can select the IP address or interface that must be used to send probe packets:
    - **None**: The source address of the probe packet is the address of the default egress interface.
    - **IPv6 Address**: The source IP address that must be used when echo request packets are sent. With this selection, the **IPv6 Address** field displays and you must enter the IPv6 address that must be used as the source.
    - **Interface**: The interface that must be used when echo request packets are sent. The Interface menu displays, and you must select the interface.
12. Click the **Apply** button.
    A traceroute request is sent to the specified IP address or host name. The results are displayed below the configurable data in the Results field.

Capture Packets

You can capture and store packets on a USB storage device.

To initiate packet capturing:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Troubleshooting > Packet Capturing**.
   The Packet Capturing page displays.
3. Next to RPCAP USB, select the **Enable** radio button.

4. From the **Capture Mode** menu, select the type of traffic that must be captured:
   - **ALL**: Capture all traffic. This is the default setting.
   - **TX**: Capture transmitted traffic only.
   - **RX**: Capture received traffic only.
5. In the **File Name** field, specify the name of the USB file.
   The file name cannot include the following symbols: '\/:*?"<>|'. You can specify up to 64 characters, excluding the extension, which is added automatically.
6. To start the packet capture process, click the **Apply** button.
   Packets are captured until you stop the process.
7. The stop the packet capture process, do the following:
   a. Next to RPCAP USB, select the **Disable** radio button.
   b. Click the **Apply** button.
      The packet capture process stops.

Perform a full memory dump

You can perform a full memory dump to retrieve the core dump for troubleshooting.

To perform a full memory dump:

1. Open a web browser and log in to the switch. For more information, see Log in to the main UI with a known IP address on page 5.
   The System Information page displays.
2. Select **Maintenance > Troubleshooting > Full Memory Dump**.
   The Full Memory Dump page displays.
3. From the **Protocol** menu, select which protocol is used to transfer the core dump file:
   - **None**: The core dump functionality is disabled.
   - **TFTP**: The results are transferred to a TFTP server.
   - **USB**: The results are transferred to a USB storage device. This is the default option.
   - **FTP**: The results are transferred to an FTP server.
4. If you select **TFTP** or **FTP** from the **Protocol** menu, additional fields display on the page, allowing you to specify the server settings:
   - **TFTP**: In the **Server Address** field, specify the IP address of the TFTP server.
   - **FTP**: Specify the following server settings:
     - **Server Address**: Specify the IP address of the FTP server.
     - **User Name**: Specify the user name for remote login to the server.
     - **Password**: Specify the password for remote login to the server.
5. In the **File Path** field, specify the path to the location to store the core dump file.

6. In the **File Name** field, specify the core dump file name.
   The default name is core.
7. To append the host name to the core dump file name, select the **Hostname** check box.
8. To append a time stamp to the core dump file name, select the **Time-stamp** check box.
   This check box is selected by default.
9. To dump the switch chip register in case of an exception, select the **Switch Register Dump** check box.
10. Select one of the following check boxes:
    - **Write Core Test**: To generate a test core dump but do not transfer it to the server or USB storage device, select the Write Core Test check box.
    - **Write Core**: To generate a core dump and transfer it to the server or USB storage device, select the Write Core check box.
11. If you select the **Write Core** check box, to save the current settings, select the **Save Current Settings** check box.
    This check box is selected by default.

    ⚠ **CAUTION:** The switch reboots after you click the **Apply** button.

12. Click the **Apply** button.
    The memory dump is sent to the specified location. The switch reboots. The process takes about 135 seconds.

Configuration Examples

This appendix contains information about how to configure the following features:

- Virtual Local Area Networks (VLANs)
- Access control lists (ACLs)
- Differentiated Services (DiffServ)
- 802.1X port access control
- Multiple Spanning Tree Protocol
- VLAN routing interfaces

Virtual Local Area Networks (VLANs)

A local area network (LAN) can generally be defined as a broadcast domain. Hubs, bridges, or switches in the same physical segment or segments connect all end node devices. End nodes can communicate with each other without the need for a router. Routers connect LANs together, routing the traffic to the appropriate port.

A virtual LAN (VLAN) is a local area network with a definition that maps workstations on some basis other than geographic location (for example, by department, type of user, or primary application). To enable traffic to flow between VLANs, traffic must go through a router, just as if the VLANs were on two separate LANs.

A VLAN is a group of computers, servers, and other network resources that behave as if they were connected to a single network segment—even though they might not be. For example, all marketing personnel might be spread throughout a building. Yet if they are all assigned to a single VLAN, they can share resources and bandwidth as if they were connected to the same segment. The resources of other departments can be invisible to the marketing VLAN members, accessible to all, or accessible only to specified individuals, depending on how the IT manager has set up the VLANs.

VLANs present a number of advantages:

- It is easy to do network segmentation. Users that communicate most frequently with each other can be grouped into common VLANs, regardless of physical location. Each group's traffic is contained largely within the VLAN, reducing extraneous traffic and improving the efficiency of the whole network.
- They are easy to manage. The addition of nodes, as well as moves and other changes, can be dealt with quickly and conveniently from a management interface rather than from the wiring closet.
- They provide increased performance. VLANs free up bandwidth by limiting node-to-node and broadcast traffic throughout the network.
- They ensure enhanced network security. VLANs create virtual boundaries that can be crossed only through a router. So standard, router-based security measures can be used to restrict access to each VLAN.

Packets received by the switch are treated in the following way:

- When an untagged packet enters a port, it is automatically tagged with the port's default VLAN ID tag number. Each port supports a default VLAN ID setting that is configurable (the default setting is 1). The default VLAN ID setting for each port can be changed on the Port PVID Configuration page. See Change the port VLAN ID settings on page 205.
- When a tagged packet enters a port, the tag for that packet is unaffected by the default VLAN ID setting. The packet proceeds to the VLAN specified by its VLAN ID tag number.
- If the port through which the packet enters is not a member of the VLAN as specified by the VLAN ID tag, the packet is dropped.
- If the port is a member of the VLAN specified by the packet's VLAN ID, the packet can be sent to other ports with the same VLAN ID.
- Packets leaving the switch are either tagged (T) or untagged (U), depending on the setting for that port's VLAN membership properties. A "U" for a port means that packets leaving the switch from that port are untagged. A "T" for a port means that packets leaving the switch from that port are tagged with the VLAN ID that is associated with the port.

The example in this section comprises numerous steps to illustrate a wide range of configurations to help provide an understanding of tagged VLANs.

VLAN configuration examples

This example demonstrates several scenarios of VLAN use and describes how the switch handles tagged and untagged traffic.

In this example, you create two new VLANs, change the port membership for default VLAN 1, and assign port members to the two new VLANs:

1. In the Basic VLAN Configuration page (see VLANs on page 195), create the following VLANs:
   - A VLAN with VLAN ID 10.
   - A VLAN with VLAN ID 20.
2. On the VLAN Membership page (see Configure membership interfaces for a VLAN on page 203) specify the VLAN membership as follows:
   - For the default VLAN with VLAN ID 1, specify the following members: port 7 (U) and port 8 (U).
   - For the VLAN with VLAN ID 10, specify the following members: port 1 (U), port 2 (U), and port 3 (T).
   - For the VLAN with VLAN ID 20, specify the following members: port 4 (U), port 5 (T), and port 6 (U).

3. On the Port PVID Configuration page (see Change the port VLAN ID settings on page 205), specify the PVID for ports g1 and g4 so that packets entering these ports are tagged with the port VLAN ID:
   - Port 1: PVID 10
   - Port 4: PVID 20
4. With the VLAN configuration that you set up, the following situations produce results as described:
   - If an untagged packet enters port 1, the switch tags it with VLAN ID 10. The packet can access port 2 and port 3. The outgoing packet is stripped of its tag to leave port 2 as an untagged packet. For port 3, the outgoing packet leaves as a tagged packet with VLAN ID 10.
   - If a tagged packet with VLAN ID 10 enters port 3, the packet can access port 1 and port 2. If the packet leaves port 1 or port 2, it is stripped of its tag to leave the switch as an untagged packet.
   - If an untagged packet enters port 4, the switch tags it with VLAN ID 20. The packet can access port 5 and port 6. The outgoing packet is stripped of its tag to become an untagged packet as it leaves port 6. For port 5, the outgoing packet leaves as a tagged packet with VLAN ID 20.

Access control lists (ACLs)

ACLs ensure that only authorized users can access specific resources while blocking off any unwarranted attempts to reach network resources.

ACLs are used to provide traffic flow control, restrict contents of routing updates, decide which types of traffic are forwarded or blocked, and provide security for the network. ACLs are normally used in firewall routers that are positioned between the internal network and an external network, such as the Internet. They can also be used on a router positioned between two parts of the network to control the traffic entering or exiting a specific part of the internal network. The added packet processing required by the ACL feature does not affect switch performance. That is, ACL processing occurs at wire speed.

Access lists are a sequential collection of permit and deny conditions. This collection of conditions, known as the filtering criteria, is applied to each packet that is processed by the switch or the router. The forwarding or dropping of a packet is based on whether or not the packet matches the specified criteria.

Traffic filtering requires the following two basic steps:

1. Create an access list definition.The access list definition includes rules that specify whether traffic matching the criteria is forwarded normally or discarded. Additionally, you can assign traffic that matches the criteria to a particular queue or redirect the traffic to a particular port. A default *deny all* rule is the last rule of every list.

2. Apply the access list to an interface in the inbound direction.

The switch allow ACLs to be bound to physical ports and LAGs. The switch supports MAC ACLs and IP ACLs.

MAC ACL sample configuration

The following example shows how to create a MAC-based ACL that permits Ethernet traffic from the Sales department on specified ports and denies all other traffic on those ports.

1. On the MAC ACL page, create an ACL with the name Sales_ACL for the Sales department of your network (see Configure a MAC ACL on page 691).By default, this ACL is bound on the inbound direction, which means that the switch examines traffic as it enters the port.
2. On the MAC Rules page, create a rule for the Sales_ACL with the following settings:
   - **Sequence Number**: 1
   - **Action**: Permit
   - **Assign Queue ID**: 0
   - **Match Every**: False
   - **CoS**: 0
   - **Destination MAC**: 01:02:1A:BC:DE:EF
   - **Destination MAC Mask**: 00:00:00:00:FF:FF
   - **EtherType**: User Value
   - **Source MAC**: 02:02:1A:BC:DE:EF
   - **Source MAC Mask**: 00:00:00:00:FF:FF
   - **VLAN ID**: 2

   For more information about MAC ACL rules, see Configure MAC ACL rules on page 693.
3. On the MAC Binding Configuration page, assign the Sales_ACL to 6, 7, and 8, and then click the **Apply** button. (See Configure MAC bindings on page 697.)You can assign an optional sequence number to indicate the order of this access list relative to other access lists if any are already assigned to this interface and direction.
4. The MAC Binding Table displays the interface and MAC ACL binding information. (See Display or delete MAC ACL bindings in the MAC binding table on page 698.)

The ACL named Sales_ACL looks for Ethernet frames with destination and source MAC addresses and MAC masks defined in the rule. Also, the frame must be tagged with VLAN ID 2, which is the Sales department VLAN. The CoS value of the frame must be 0, which is the default value for Ethernet frames. Frames that match this criteria are permitted on interfaces 6, 7, and 8 and are assigned to the hardware egress queue 0, which is the default queue. All other traffic is explicitly denied on these interfaces. To allow additional traffic to enter these ports, you must add a new Permit rule with the desired match criteria and bind the rule to interfaces 6, 7, and 8.

Basic IP ACL sample configuration

The following example shows how to create an IP-based ACL that prevents any IP traffic from the Finance department from being allowed on the ports that are associated with other departments. Traffic from the Finance department is identified by each packet's network IP address.

1. On the IP ACL page, create a new IP ACL with an IP ACL ID of 1. (See Configure a basic or extended IPv4 ACL on page 699.)
2. On the IP Rules page, create a rule for IP ACL 1 with the following settings:
   - **Sequence Number**: 1
   - **Action**: Deny
   - **Assign Queue ID**: 0 (optional: 0 is the default value)
   - **Match Every**: False
   - **Source IP Address**: 192.168.187.0
   - **Source IP Mask**: 255.255.0

   For additional information about IP ACL rules, see Configure rules for a basic IP ACL on page 701.
3. Click the **Add** button.
4. On the IP Rules page, create a second rule for IP ACL 1 with the following settings:
   - **Sequence Number**: 2
   - **Action**: Permit
   - **Match Every**: True
5. Click the **Add** button.
6. On the IP Binding Configuration page, assign ACL ID 1 to ports 2, 3, and 4, and assign a sequence number of 1. (See Configure IP ACL interface bindings on page 720.)
   By default, this IP ACL is bound on the inbound direction, so it examines traffic as it enters the switch.
7. Click the **Apply** button.
8. On the IP Binding Table page, you can view the interfaces and IP ACL binding information. (See Display or delete IP ACL bindings in the IP ACL binding table on page 722.)

   The IP ACL in this example matches all packets with the source IP address and subnet mask of the Finance department's network and deny it on interfaces 2, 3, and 4 of the switch. The second rule permits all non-Finance traffic on the ports. The second rule is required because an explicit *deny all* rule exists as the lowest priority rule.

Differentiated Services (DiffServ)

Standard IP-based networks are designed to provide *best effort* data delivery service. *Best effort* service implies that the network delivers the data in a timely fashion, although there is no guarantee that it does. During times of congestion, packets might be delayed, sent sporadically, or dropped. For typical Internet applications, such as email and file transfer, a slight degradation in service is acceptable and in many cases unnoticeable. However, any degradation of service can negatively affect applications with strict timing requirements, such as voice or multimedia. Quality of Service (QoS) can provide consistent, predictable data delivery by distinguishing between packets with strict timing requirements and those that are more tolerant of delay. Packets with strict timing requirements are given special treatment in a QoS-capable network. With this in mind, all elements of the network must be QoS capable. If one node cannot meet the necessary timing requirements, this creates a deficiency in the network path and the performance of the entire packet flow is compromised.

There are two basic types of QoS:

- **Integrated Services**: Network resources are apportioned based on request and are reserved (resource reservation) according to network management policy (RSVP, for example).
- **Differentiated Services**: Network resources are apportioned based on traffic classification and priority, giving preferential treatment to data with strict timing requirements.

The switch supports DiffServ.

The DiffServ feature contains a number of conceptual QoS building blocks that you can use to construct a differentiated service network. Use these same blocks in different ways to build other types of QoS architectures.

You must configure three key QoS building blocks for DiffServ:

- Class
- Policy
- Service (the assignment of a policy to a directional interface)

Class

You can classify incoming packets at Layers 2, 3 and 4 by inspecting the following information for a packet:

- Source/destination MAC address
- EtherType
- Class of Service (802.1p priority) value (first/only VLAN tag)
- VLAN ID range (first/only VLAN tag)
- Secondary 802.1p priority value (second/inner VLAN tag)
- Secondary VLAN ID range (second/inner VLAN tag)
- IP Service Type octet (also known as: ToS bits, Precedence value, DSCP value)
- Layer 4 protocol (TCP, UDP and so on)

- Layer 4 source/destination ports
- Source/destination IP address

From a DiffServ point of view, there are two types of classes:

- DiffServ traffic classes
- DiffServ service levels/forwarding classes

DiffServ traffic classes

With DiffServ, you define which traffic classes to track on an ingress interface. You can define simple behavior aggregate (BA) classifiers (DSCP) and a wide variety of multi-field (MF) classifiers:

- Layer 2; Layers 3, 4 (IP only)
- Protocol-based
- Address-based

You can combine these classifiers with logical AND or OR operations to build complex MF-classifiers (by specifying a class type of *all* or *any*, respectively). That is, within a single class, multiple match criteria are grouped together as an AND expression or a sequential OR expression, depending on the defined class type. Only classes of the same type can be nested; class nesting does not allow for the negation (*exclude* option) of the referenced class.
To configure DiffServ, you must define service levels, namely the forwarding classes/per-hop behaviors (PHBs) identified by a DSCP value, on the egress interface. You define these service levels by configuring BA classes for each.

Create policies

Use DiffServ policies to associate a collection of classes that you configure with one or more QoS policy statements. The result of this association is referred to as a policy.
From a DiffServ perspective, there are two types of policies:

- **Traffic Conditioning Policy**: A policy applied to a DiffServ traffic class
- **Service Provisioning Policy**: A policy applied to a DiffServ service level

You must manually configure the various statements and rules used in the traffic conditioning and service provisioning policies to achieve the desired Traffic Conditioning Specification (TCS) and the Service Level Specification (SLS) operation, respectively.

Traffic conditioning policy

Traffic conditioning pertains to actions performed on incoming traffic. Several distinct QoS actions are associated with traffic conditioning:

- **Dropping**: Drop a packet upon arrival. This is useful for emulating access control list operation using DiffServ, especially when DiffServ and ACL cannot coexist on the same interface.

- **Marking IP DSCP or IP precedence**: Marking/re-marking the DiffServ code point in a packet with the DSCP value representing the service level associated with a particular DiffServ traffic class. Alternatively, the IP Precedence value of the packet can be marked/re-marked.
- **Marking CoS (802.1p)**: Sets the three-bit priority field in the first/only 802.1p header to a specified value when packets are transmitted for the traffic class. An 802.1p header is inserted if it does not already exist. This is useful for assigning a Layer 2 priority level based on a DiffServ forwarding class (such as the DSCP or IP precedence value) definition to convey some QoS characteristics to downstream switches that do not routinely look at the DSCP value in the IP header.
- **Policing**: A method of constraining incoming traffic associated with a particular class so that it conforms to the terms of the TCS. Special treatment can be applied to out-of-profile packets that are either in excess of the conformance specification or are non-conformant. The DiffServ feature supports the following types of traffic policing treatments (actions):
  - **Send**: The packet is forwarded without DiffServ modification.
  - **Drop**: The packet is dropped.
  - **Mark CoS**: The 802.1p user priority bits are (re)marked and forwarded.
  - **Mark DSCP**: The packet DSCP is (re)marked and forwarded.
  - **Mark IP precedence**: The packet IP Precedence is (re)marked and forwarded.
- **Color mode awareness**: Policing in the DiffServ feature uses either *color blind* or *color aware* mode. Color blind mode ignores the coloration (marking) of the incoming packet. Color aware mode takes into consideration the current packet marking when determining the policing outcome. An auxiliary traffic class is used in conjunction with the policing definition to specify a value for one of the 802.1p, secondary 802.1p, IP DSCP, or IP Precedence fields designating the incoming color value to be used as the conforming color. As an option, you can also specify the color of exceeding traffic.
- **Assigning QoS queue**: Directs a traffic stream to the specified QoS queue. This allows a traffic classifier to specify which one of the supported hardware queues is used for handling packets belonging to the class.
- **Redirecting**. Forces a classified traffic stream to a specified egress port (physical or LAG). This can occur in addition to any marking or policing action. It can also be specified along with a QoS queue assignment.
- **Mirroring**. Copies a classified traffic stream to a specified egress port (physical or LAG). This can occur in addition to any marking or policing action. It can also be specified along with a QoS queue assignment.
- **Counting**: Updating octet and packet statistics to keep track of data handling along traffic paths within DiffServ. In this DiffServ feature, counters are not explicitly configured by the user, but are designed into the system based on the DiffServ policy being created. For more information, see Port and EAP packet statistics on page 725.

DiffServ example configuration

To create a DiffServ Class/Policy and attach it to a switch interface, follow these steps:

1. On the QoS Class Configuration page, create a new class with the following settings:
   - **Class Name**: Class1
   - **Class Type**: All

   For more information about this page, see Configure a DiffServ class on page 532.
2. Click the **Class1** hyperlink to display the DiffServ Class Configuration page for this class.
3. Configure the following settings for Class1:
   - **Protocol Type**: UDP
   - **Source IP Address**: 192.12.1.0.
   - **Source Mask**: 255.255.255.0.
   - **Source L4 Port**: Other, and enter 4567 as the source port value.
   - **Destination IP Address**: 192.12.2.0.
   - **Destination Mask**: 255.255.255.0.
   - **Destination L4 Port**: Other, and enter 4568 as the destination port value.

   For more information about this page, see Configure a DiffServ class on page 532.
4. Click the **Apply** button.
5. On the Policy Configuration page, create a new policy with the following settings:
   - **Policy Selector**: Policy1
   - **Member Class**: Class1

   For more information about this page, see Configure a DiffServ policy on page 543.
6. Click the **Add** button.
   The policy is added.
7. Click the **Policy1** hyperlink to display the Policy Class Configuration page for this policy.
8. Configure the Policy attributes as follows:
   - **Assign Queue**: 3
   - **Policy Attribute**: Simple Policy
   - **Color Mode**: Color Blind
   - **Committed Rate**: 1000000 Kbps
   - **Committed Burst Size**: 128 KB
   - **Confirm Action**: Send
   - **Violate Action**: Drop

   For more information about this page, see Configure a DiffServ policy on page 543.
9. On the Service Configuration page, select the check box next to interfaces 7 and 8 to attach the policy to these interfaces, and then click the **Apply** button. (See Configure the DiffServ service interface on page 550.)

All UDP packet flows destined to the 192.12.2.0 network with an IP source address from the 192.12.1.0 network that include a Layer 4 Source port of 4567 and Destination port of 4568 from this switch on ports 7 and 8 are assigned to hardware queue 3.
On this network, traffic from streaming applications uses UDP port 4567 as the source and 4568 as the destination. This real-time traffic is time sensitive, so it is assigned to a high-priority hardware queue. By default, data traffic uses hardware queue 0, which is designated as a best-effort queue.
Also the *confirmed action* on this flow is to send the packets with a committed rate of 1000000 Kbps and burst size of 128 KB. Packets that violate the committed rate and burst size are dropped.

802.1X port access control

Local Area Networks (LANs) are often deployed in environments that permit unauthorized devices to be physically attached to the LAN infrastructure, or permit unauthorized users to attempt to access the LAN through equipment already attached. In such environments you might want to restrict access to the services offered by the LAN to those users and devices that are permitted to use those services.
Port-based network access control makes use of the physical characteristics of LAN infrastructures to provide a means of authenticating and authorizing devices attached to a LAN port with point-to-point connection characteristics. If the authentication and authorization process fails, access control prevents access to that port. In this context, a port is a single point of attachment to the LAN, such as a port of a MAC bridge and an association between stations or access points in IEEE 802.11 wireless LANs.
The IEEE 802.11 standard describes an architectural framework within which authentication and consequent actions take place. It also establishes the requirements for a protocol between the authenticator (the system that passes an authentication request to the authentication server) and the supplicant (the system that requests authentication), as well as between the authenticator and the authentication server.
The switch support a guest VLAN, which allows unauthenticated users limited access to network resources.

NOTE: You can use QoS features to provide rate limiting on the guest VLAN to limit the network resources the guest VLAN provides.

Another 802.1X feature is the ability to configure a port to enable or disable EAPoL packet forwarding support. You can disable or enable the forwarding of EAPoL when 802.1X is disabled on the device.
The ports of an 802.1X authenticator switch provide the means by which it can offer services to other systems reachable through the LAN. Port-based network access control allows the

operation of a switch's ports to be controlled to ensure that access to its services is permitted only by systems that are authorized to do so.

Port access control provides a means of preventing unauthorized access by supplicants to the services offered by a system. Control over the access to a switch and the LAN to which it is connected can be desirable if you restrict access to publicly accessible bridge ports or departmental LANs.

Access control is achieved by enforcing authentication of supplicants that are attached to an authenticator's controlled ports. The result of the authentication process determines whether the supplicant is authorized to access services on that controlled port.

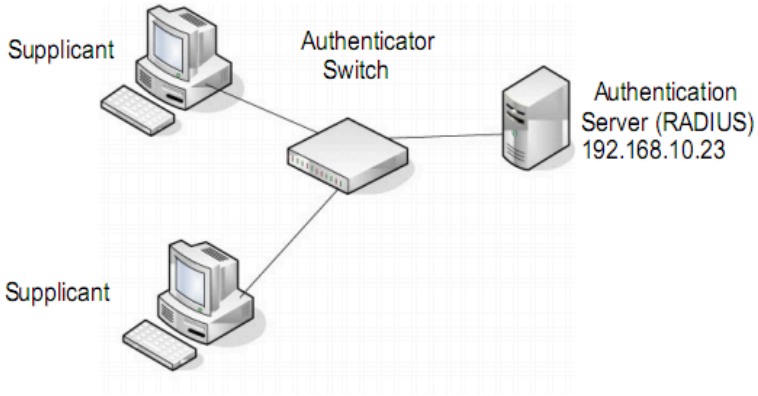
A Port Access Entity (PAE) is able to adopt one of two distinct roles within an access control interaction:

1. **Authenticator**: A port that enforces authentication before allowing access to services available through that port.
2. **Supplicant**: A port that attempts to access services offered by the authenticator.

In addition, an authentication server is required. This is a device that performs the authentication function necessary to check the credentials of the supplicant on behalf of the authenticator. To complete an authentication exchange, an authenticator, supplicant, and authentication server are required.

The switch support the authenticator role only, in which the PAE is responsible for communicating with the supplicant. The authenticator PAE is also responsible for submitting the information received from the supplicant to the authentication server for the credentials to be checked, which determines the authorization state of the port. The authenticator PAE controls the authorized/ unauthorized state of the controlled port depending on the outcome of the RADIUS-based authentication process.

Figure 1. 802.1X authentication roles



802.1X example configuration

This example shows how to configure the switch so that 802.1X-based authentication is required on the ports in a corporate conference room (5 through 8). These ports are available to visitors and must be authenticated before access is granted to the network. The authentication is handled by an external RADIUS server. When the visitor is successfully authenticated, traffic is automatically assigned to the guest VLAN. This example assumes that a VLAN was configured with a VLAN ID of 150 and VLAN name of Guest.

1. On the Port Authentication page, select ports **5**, **6**, **7**, and **8**.
2. From the **Port Control** menu, select **Unauthorized**.
   The selection from the **Port Control** menu for all other ports on which authentication is not needed must be **Authorized**. When the selection from the **Port Control** menu is **Authorized**, the port is unconditionally put in a force-authorized state and does not require any authentication. When the selection from the **Port Control** menu is **Auto**, the authenticator PAE sets the controlled port mode.
3. In the **Guest VLAN** field for ports 5 through 8, enter **150** to assign these ports to the guest VLAN.
   You can configure additional settings to control access to the network through the ports. See Configure a port security interface on page 620 for information about the settings.
4. Click the **Apply** button.

5. On the 802.1X Configuration page, set the Port Based Authentication State and Guest VLAN mode to **Enable**, and then the **Apply** button (See Configure the global port security mode on page 619).
   This example uses the default values for the port authentication settings, but you can configure several additional settings. For example, the **EAPOL Flood Mode** field allows you to enable the forwarding of EAPoL frames when 802.1X is disabled on the device.
6. On the RADIUS Server Configuration page, configure a RADIUS server with the following settings:
   - **Server Address**: 192.168.10.23
   - **Secret Configured**: Yes
   - **Secret**: secret123
   - **Active**: Primary

   For more information, see RADIUS servers on page 562.
7. Click the **Add** button.
8. On the Authentication List page, configure the default List to use RADIUS as the first authentication method (See Configure a login authentication list on page 574).
   This example enables 802.1X-based port security on the switch and prompts the hosts connected on ports 5 through 8 for an 802.1X-based authentication. The switch passes the authentication information to the configured RADIUS server.

Multiple Spanning Tree Protocol

Spanning Tree Protocol (STP) runs on bridged networks to help eliminate loops. If a bridge loop occurs, the network can become flooded with traffic. IEEE 802.1s Multiple Spanning Tree Protocol (MSTP) supports multiple instances of spanning tree to efficiently channel VLAN traffic over different interfaces. Each instance of the spanning tree behaves in the manner specified in IEEE 802.1w, Rapid Spanning Tree, with slight modifications in the working but not the end effect (chief among the effects is the rapid transitioning of the port to the forwarding state).
The difference between RSTP and traditional STP (IEEE 802.1D) is the ability to configure and recognize full-duplex connectivity and ports that are connected to end stations, resulting in rapid transitioning of the port to the Forwarding state and the suppression of Topology Change Notification. These features are represented by the parameters pointtopoint and edgeport.
MSTP is compatible with both RSTP and STP. It behaves in a way that is appropriate for STP and RSTP bridges.
An MSTP bridge can be configured to behave entirely as an RSTP bridge or an STP bridge. So, an IEEE 802.1s bridge inherently also supports IEEE 802.1w and IEEE 802.1D.
The MSTP algorithm and protocol provide simple and full connectivity for frames assigned to any VLAN throughout a bridged LAN comprising arbitrarily interconnected networking devices, each operating with MSTP, STP, or RSTP. MSTP allows frames assigned to different VLANs to follow

separate paths, each based on an independent Multiple Spanning Tree Instance (MSTI), within Multiple Spanning Tree (MST) regions composed of LANs and or MSTP bridges. These regions and the other bridges and LANs are connected into a single Common Spanning Tree (CST). (IEEE DRAFT P802.1s/D13)

MSTP connects all bridges and LANs with a single Common and Internal Spanning Tree (CIST). The CIST supports the automatic determination of each MST region, choosing its maximum possible extent. The connectivity calculated for the CIST provides the CST for interconnecting these regions, and an Internal Spanning Tree (IST) within each region.

MSTP ensures that frames with a VLAN ID are assigned to one and only one of the MSTIs or the IST within the region, that the assignment is consistent among all the networking devices in the region, and that the stable connectivity of each MSTI and IST at the boundary of the region matches that of the CST.

MSTP ensures the following:

- frames with a VLAN ID are assigned to one and only one of the MSTIs or the IST within the region
- the assignment is consistent among all the networking devices in the region
- the stable connectivity of each MSTI and IST at the boundary of the region matches that of the CST

The stable active topology of the bridged LAN with respect to frames consistently classified as belonging to any VLAN thus simply and fully connects all LANs and networking devices throughout the network, though frames belonging to different VLANs can take different paths within any region, per IEEE DRAFT P802.1s/D13.

All bridges, whether they use STP, RSTP, or MSTP, send information in configuration messages through Bridge Protocol Data Units (BPDUs) to assign port roles that determine each port's participation in a fully and simply connected active topology based on one or more spanning trees. The information communicated is known as the spanning tree priority vector. The BPDU structure for each of these different protocols is different. An MSTP bridge transmits the appropriate BPDU depending on the received type of BPDU from a particular port.

An MST region comprises of one or more MSTP bridges with the same MST configuration identifier, using the same MSTIs, and without any bridges attached that cannot receive and transmit MSTP BPDUs. The MST configuration identifier includes the following components:

- Configuration identifier format selector
- Configuration name
- Configuration revision level
- Configuration digest: 16-byte signature of type HMAC-MD5 created from the MST Configuration Table (a VLAN ID to MSTID mapping)

Because multiple instances of spanning tree exist, an MSTP state is maintained on a per-port, per-instance basis (or on a per-port, per-VLAN basis, as any VLAN can be in one and only one MSTI or CIST). For example, port A can be forwarding for instance 1 while discarding for instance 2. The port states changed since the IEEE 802.1D specification.

To support multiple spanning trees, configure an MSTP bridge with an unambiguous assignment of VLAN IDs (VIDs) to spanning trees. For such a configuration, ensure the following:

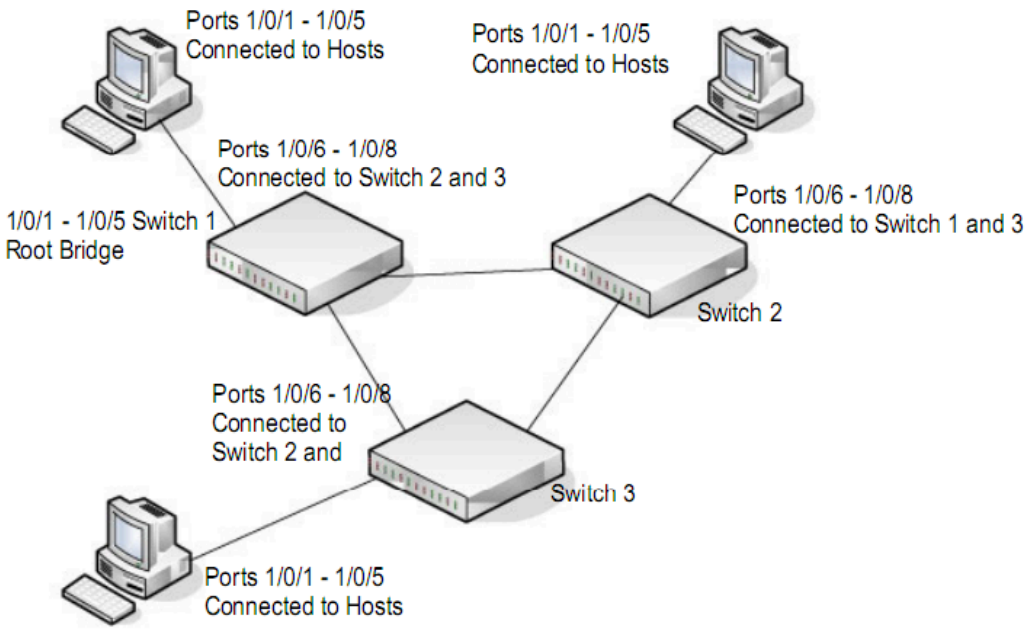
- The allocation of VIDs to FIDs is unambiguous.
- Each FID that is supported by the bridge is allocated to exactly one spanning tree instance.

The combination of VID to FID and then FID to MSTI allocation defines a mapping of VIDs to spanning tree instances, represented by the MST Configuration Table.
With this allocation we ensure that every VLAN is assigned to one and only one MSTI. The CIST is also an instance of spanning tree with an MSTID of 0.
VIDs might be not be allocated to an instance, but every VLAN must be allocated to one of the other instances of spanning tree.
The portion of the active topology of the network that connects any two bridges in the same MST region traverses only MST bridges and LANs in that region, and never bridges of any kind outside the region. In other words, connectivity within the region is independent of external connectivity.

MSTP example configuration

This example shows how to create an MSTP instance from the switch. The example network includes three different switches that serve different locations in the network. In he following figure, ports 1/0/1–1/0/5 are connected to host stations, so those links are not subject to network loops. Ports 1/0/6–1/0/8 are connected across switches 1, 2, and 3.

Figure 2. MSTP sample configuration



Perform the following procedures on each switch to configure MSTP:

1. Use the VLAN Configuration page to create VLANs 300 and 500 (see Change the internal VLAN allocation settings on page 198).
2. On the VLAN Membership page, include ports 1 through 8 as tagged (T) or untagged (U) members of VLAN 300 and VLAN 500 (see Change the internal VLAN allocation settings on page 198).
3. On the STP Configuration page, enable the Spanning Tree State option (see Configure the CST interface settings on page 235).
   Use the default values for the rest of the STP configuration settings. By default, the STP operation mode is MSTP and the configuration name is the switch MAC address.
4. On the CST Configuration page, set the bridge priority value for each of the three switches to force Switch 1 to be the root bridge:
   - **Switch 1**: 4096
   - **Switch 2**: 12288
   - **Switch 3**: 20480

   NOTE: Bridge priority values are multiples of 4096.

   If you do not specify a root bridge and all switches are assigned the same bridge priority value, the switch with the lowest MAC address is elected as the root bridge (see Configure the CST settings and display the CST status on page 232).
5. On the CST Port Configuration page, select ports 1 through 8 and select **Enable** from the **STP Status** menu (see Configure the CST interface settings on page 235).
6. Click the **Apply** button.
7. Select ports 1 through 5 (edge ports), and select **Enable** from the **Fast Link** menu.
   Because the edge ports are not at risk for network loops, ports with Fast Link enabled transition directly to the forwarding state.
8. Click the **Apply** button.
   You can use the CST Port Status page to view spanning tree information about each port.
9. On the MST Configuration page (Manage MST instances on page 240), create an MST instance with the following settings:
   - **MST ID**: 1
   - **Priority**: Use the default (32768)
   - **VLAN ID**: 300
10. Click the **Add** button.

11. Create a second MST instance with the following settings:
    - **MST ID**: 2
    - **Priority**: 49152
    - **VLAN ID**: 500
12. Click the **Add** button.

    In this example, assume that Switch 1 became the root bridge for the MST instance 1, and Switch 2 became the root bridge for MST instance 2. Switch 3 supports hosts in the sales department (ports 1, 2, and 3) and in the HR department (ports 4 and 5). Switches 1 and 2 also include hosts in the sales and HR departments. The hosts connected from Switch 2 use VLAN 500, MST instance 2 to communicate with the hosts on Switch 3 directly. Likewise, hosts of Switch 1 use VLAN 300, MST instance 1 to communicate with the hosts on Switch 3 directly.
    The hosts use different instances of MSTP to effectively use the links across the switch. The same concept can be extended to other switches and more instances of MSTP.

VLAN routing interfaces

VLANs divide broadcast domains in a LAN environment. When hosts in one VLAN must communicate with hosts in another VLAN, the traffic must be routed between them. This is known as inter-VLAN routing. On the switch, it is accomplished by creating Layer 3 interfaces (switch virtual interfaces [SVI]).
When a port is enabled for bridging (the default) rather than routing, all normal bridge processing is performed for an inbound packet, which is then associated with a VLAN. Its MAC destination address and VLAN ID are used to search the MAC address table. If routing is enabled for the VLAN, and the MAC destination address of an inbound unicast packet is that of the internal bridge-router interface, the packet is routed. An inbound multicast packet is forwarded to all ports in the VLAN, plus the internal bridge-router interface, if it was received on a routed VLAN.
Because a port can be configured to belong to more than one VLAN, VLAN routing might be enabled for all of the VLANs on the port, or for a subset. VLAN routing can be used to allow more than one physical port to reside on the same subnet. It could also be used when a VLAN spans multiple physical networks, or when additional segmentation or security is required. A port can be either a VLAN port or a router port, but not both. However, a VLAN port can be part of a VLAN that is itself a router port.
Complete these steps to configure a switch to perform interVLAN routing for IPv4:

1. Use the IP Configuration page to enable IPv4 routing on the switch.
   By default, IPv4 routing is enabled. For more information about this step, see Manage the global IPv4 routing settings on page 354.

2. Determine the IP addresses that you want to assign to the VLAN interface on the switch.
   For the switch to be able to route between the VLANs, the VLAN interfaces must be configured with an IP address. When the switch receives a packet destined for another subnet or VLAN, the switch looks at the routing table to determine where to forward the packet. The packet is then passed to the VLAN interface of the destination. It is then sent to the port where the end device is attached.
3. Use the VLAN Static Routing Wizard page to create a routing VLAN, configure the IP address and subnet mask, and add the member ports.
   For more information about this step, see Create a routing VLAN with the VLAN static routing wizard on page 379.

Software Default Settings and Hardware Specifications

This appendix contains the following sections:

- Access default settings for the switch main UI
- System features default settings
- Switching features default settings
- Routing, OSPF, OSPFv3, and multicast features default settings
- QoS features default settings
- Security features default settings
- Monitoring features default settings
- General hardware technical specifications
- Model-specific hardware technical specifications

NOTE: For more information about the switch specifications and capabilities, including the maximum settings for many features, see the data sheet, which you can download by visiting netgear.com/support/download/.

Access default settings for the switch main UI

The following table describes the default settings for access to the switch main UI. (View-only settings are not included in the table but might be included in the data sheet, which you can download by visiting netgear.com/support/download/.)

Table 205.Default settings for access to the switch main UI

Feature	Default
IP address for management VLAN	169.254.100.100
Service port IP address	192.168.0.239
Subnet mask	255.255.0.0
Default gateway	0.0.0.0
Management VLAN ID	1
Service port protocol	DHCP client enabled
Admin user	User name: admin (read/write access) Password: No password (that is, the password is blank), but upon first login, the admin user must specify a password.

Feature	Default
	Encryption type: SHA512 Multifactor authentication mode: Disabled
Guest user	User name: guest (read-only access) Password: No password (that is, the password is blank), but upon first login, the guest user must specify a password. Encryption type: SHA512 Multifactor authentication mode: Disabled
Minimum length for admin and guest passwords	Eight characters
IPv6 management mode	Enabled
IPv6 address auto configuration	Disabled

System features default settings

The following table describes the default settings for the system features that you can configure. (View-only settings are not included in the table but might be included in the data sheet, which you can download by visiting netgear.com/support/download/.)

Table 206.System features default settings

Feature Name/Setting	Default
Time, Global	
Clock source	NTP
Time, NTP, Global	
NTP authentication mode	Enabled
NTP broadcast client mode	Disabled
Broadcast delay	3000
Source interface	VLAN 1
Offset hours	0
Offset minutes	0
Time, NTP Server	
Server type	DNS (for each default server)
Address	time-a.netgear.com

Feature Name/Setting	Default
	time-c.netgear.com
Version	4 (for each default server)
Key ID	0 (for each default server)
Minimum poll interval	6 (for each default server)
Maximum poll interval	10 (for each default server)
Prefer	No
Burst	Yes
Iburst	Yes
Time, Daylight Saving	
Daylight saving	Disabled
DNS	
DNS status	Enabled
Retry number	2
Response time-out	3
Source interface	VLAN 1
DNS server	8.8.8.8
SDM template preference	
SDM template preference	IPv4-Basic
Green Ethernet, Global	
Auto Power Down mode	Disabled
EEE mode	Disabled
Green Ethernet, Interface	
Auto Power Down mode	Disabled
EEE mode	Disabled
Energy Detect admin mode	Disabled
EEE admin mode	Disabled
EEE transmit idle time	600
EEE transmit wake time	17

Feature Name/Setting	Default
Green Ethernet LPI History	
Sampling interval	3600
Maximum samples to keep	168
Bonjour	
Admin mode	Enabled
Precision Time Protocol, Global	
Admin mode	Disabled
Precision Time Protocol, Interface	
Configured mode	Disabled
TFTP Server	
Admin mode	Disabled
DHCP Server No DHCP servers configured	
Admin mode	Disabled
Ping packet count	2
Conflict logging mode	Enabled
Bootp automatic mode	Disabled
DHCP Relay No DHCP relays configured	
Maximum hop count	4
Admin mode	Disabled
Minimum wait time	0
Circuit ID option mode	None
Server override mode	Disabled
DHCP L2 Relay, Global No DHCP L2 relays configured	
Admin mode	Disabled
DHCP L2 Relay, Interface	
Admin mode	Disabled

Feature Name/Setting	Default
82 Option trust mode	Disabled
No Option82 mode	Drop
UDP Relay No UDP relays configured	
Admin mode	Disabled
DHCPv6 Server, Global No DHCPv6 servers configured	
Admin mode	Disabled
DHCPv6 Server, Interface	
Admin mode	Disabled
DHCPv6 Relay No DHCPv6 relays configured	
Admin mode	Disabled
PoE Configuration (PoE+ and PoE++ models only)	
System usage threshold	90 percent
Power management mode	Dynamic
Traps	Enabled
PoE Port Configuration (PoE+ and PoE++ models only)	
Port power	Enabled
Port priority	Low
Power mode	802.3at for PoE+ models 802.3bt for PoE++ models
Power limit type	Class
Power limit (W)	90.90
Detection type	4ptdot3af
Timer schedule	None
SNMPv1/v2	
Community configuration	None
Trap configuration	None

Feature Name/Setting	Default
SNMPv1/v2 Trap Flags	
Authentication	Enabled
Link up/down	Enabled
Multiple users	Enabled
Spanning tree	Enabled
ACL	Disabled
Captive portal	Disabled
DVMRP	Disabled
PIM	Disabled
PoE (for PoE+ and PoE++ models only)	Enabled
OSPFv2 Traps, Errors	All disabled
OSPFv2 Traps, LSA	All disabled
OSPFv2 Traps, Overflow	All disabled
OSPFv2 Traps, Retransmit	All disabled
OSPFv2 Traps, State-change	All disabled
OSPFv3 Traps, Errors	All disabled
OSPFv3 Traps, LSA	All disabled
OSPFv3 Traps, Overflow	All disabled
OSPFv3 Traps, Retransmit	All disabled
OSPFv3 Traps, State-change	All disabled
Power Supply Module state	Enabled
Fan status	Enabled
Temperature status	Enabled
VRRP	Enabled
MAC notification	Disabled
SNMPv3	
User name	admin
SNMPv3 access mode	Read/Write

Feature Name/Setting	Default
Authentication protocol	SHA512
Encryption protocol	None
Trap configuration	None
LLDP, Global	
TLV advertised interval	30
Hold multiplier	4
Reinitializing delay	2
Transmit delay	5
LLDP, Interface	
Transmit	Enabled
Receive	Enabled
Notify	Enabled
Optional TLVs (port description, system name, system description, system capabilities)	All enabled
Management information	Enabled
LLDP-MED, Global	
Fast start repeat count	3
LLDP-MED, Interface	
MED status	Enabled
Notification status	Enabled
Transmit type length values (MED capabilities, network policy, location identification, extended MDI-PSE, inventory information)	All enabled
Link Dependencies None configured	
ISDP, Global	
Admin mode	Enabled
Timer	30
Hold time	180

Feature Name/Setting	Default
Version 2 advertisements	Enabled
ISDP, Interface	
Admin mode	Enabled
Timer Schedules No timer schedules configured	

Switching features default settings

The following table describes the default settings for the switching features that you can configure.
(View-only settings are not included in the table but might be included in the data sheet, which you can download by visiting netgear.com/support/download/.)

Table 207.Switching features default settings

Feature Name/Setting	Default
VLANs, Global	
Internal VLAN allocation base	4093
Internal VLAN allocation policy	Ascending
Default VLAN ID	1 (All ports are members)
Default VLAN name	default
VLANs Trunking, Global	
Admin mode	Enabled
Switch Ports, Interface	
Switch port mode	General
Access VLAN ID	1
Native VLAN ID	1
Trunk allowed VLANs	1–4093
Port PVIDs, Interface	
PVID	1
VLAN member	1
VLAN tag	None

Feature Name/Setting	Default
Acceptable Frame	Admit All
Ingress filtering	Disabled
Port priority	0
DVLAN, Interface	
Admin mode	Disabled
Voice VLAN, Global	
Admin mode	Disabled
Voice VLAN, Interface	
Interface mode	Disabled
Value	0
Cos override mode	Disabled
Authentication mode	Enabled
DSCP value	0
GARP, Global	
GVRP mode	Disabled
GMRP mode	Disabled
GARP, Interface	
GVRP mode	Disabled
GMRP mode	Disabled
Join timer	20
Leave timer	60
Leave all timer	1000
Auto-VoIP Protocol-Based, Global	
Prioritization type	Traffic Class
Class value	6
Auto-VoIP Protocol-Based, Interface	
Auto-VoIP mode	Disabled
Auto-VoIP OUI-Based, Global	

Feature Name/Setting	Default
Auto-VoIP VLAN ID	None (0)
OUI-based priority	7
OUI table	00:01:E3 SIEMENS 00:03:6B CISCO1 00:12:43 CISCO2 00:0F:E2 H3C 00:60:B9 NITSUKO 00:D0:1E PINTEL 00:E0:75 VERILINK 00:E0:BB 3COM 00:04:0D AVAYA1 00:1B:4F AVAYA2 00:04:13 SNOM
Auto-VoIP OUI-based, Interface	
Auto-VoIP mode	Disabled
iSCSI	
iSCSI status	Disabled
QoS profile	VLAN Priority Tag
VLAN priority tag	5
DSCP	46
Remark	Enabled
iSCSI aging time	10
iSCSI targets, TCP ports	860 3260
STP, Global	
Spanning tree admin mode	Enabled
Force protocol version	EEE 802.1w
Configuration name	MAC address
Configuration revision level	0
Forwarding of BPDUs while STP is disabled	Disabled
BPDU guard	Disabled

Feature Name/Setting	Default
BPDU filter	Disabled
Fast backbone	Disabled
Fast uplink	Disabled
Max update rate	150
CST, Global	
Bridge priority	32768
Bridge maximum age	20
Bridge forward delay	15
Spanning tree maximum hops	20
Spanning tree Tx hold count	6
CST, Interface	
Port priority	128
Admin edge port	Disabled
Port path cost	0
Auto calculated port path cost	Enabled
External port path cost	0
Auto calculated external port path cost	Enabled
BPDU filter	Disabled
BPDU forwarding	Disabled
Auto edge	Enabled
Root guard	Disabled
Loop guard	Disabled
TCN guard	Disabled
Port mode	Enabled
MST, Global	
MST ID	0
Priority	32768
VLAN ID	1

Feature Name/Setting	Default
PVST, VLAN	
VLAN ID	1
Root	None
Hello time	2
Forward time	15
Maximum age	20
Priority	32768
PVST, Interface	
Priority	128
Cost	0
Auto calculated cost	Enabled
IGMP Snooping, Global	
Admin mode	Enabled
Validate IGMP IP header	Enabled
Proxy Querier mode	Enabled
Report Flood mode	Enabled
Exclude Mrouter Interface mode	Enabled
Fast Leave Auto-Assignment mode	Enabled
IGMP Plus mode	Enabled
IGMP Snooping, Interface	
Admin mode	Disabled
Membership interval	600
Maximum response time	120
Expiration time	300
Fast leave	Disabled
Proxy querier	Enabled
IGMP Snooping, VLAN	
VLAN ID	1

Feature Name/Setting	Default
Admin mode	Enabled
Fast leave	Enabled
Membership interval	600
Maximum response time	120
Multicast router expiry time	300
Report suppression	Disabled
Proxy querier	Enabled
Report Flood mode	Enabled
Exclude Mrouter Interface mode	Enabled
IGMP Plus mode	Enabled
IGMP Snooping, Multicast Router, Interface	
Multicast router	Disabled
IGMP Snooping, Querier, Global	
Querier admin mode	Enabled
Snooping querier address	None
IGMP version	2
Query interval	60
Querier expiry interval	180
IGMP Snooping, Querier, VLAN	
VLAN ID	1
Querier Election Participate mode	Enabled
Querier VLAN address	None
MLD Snooping, Global	
MLD snooping admin mode	Enabled
Proxy Querier mode	Enabled
Report Flood mode	Enabled
Exclude Mrouter Interface mode	Enabled
MLD Plus mode	Enabled

Feature Name/Setting	Default
MLD Snooping, Interface	
Admin mode	Disabled
Membership interval	260
Maximum response time	10
Expiration time	300
Fast leave	Disabled
Proxy querier	Enabled
MLD Snooping, VLAN	
VLAN ID	1
Fast leave	Enabled
Membership interval	260
Maximum response time	10
Multicast router expiry time	300
Proxy Querier mode	Enabled
Exclude Mrouter Interface mode	Enabled
MLD Plus mode	Enabled
MLD Snooping, Multicast Router, Interface	
Multicast router	Disabled
MLD Snooping, Querier, Global	
Querier admin mode	Disabled
Querier address	None
Query interval	60
Querier expiry interval	60
MVR, Global	
MVR running	Disabled
MVR multicast VLAN	1
MVR global query response time	5
MVR mode	Compatible

Feature Name/Setting	Default
MVR, Interface	
Admin mode	Disabled
Type	None
Immediate leave	Disabled
Address Table	
Address aging timeout	3000
Ports	
STP mode	Enabled
Admin mode	Enabled
LACP mode	Enabled
Autonegotiation	Enabled
Speed	Auto
Duplex mode	Auto
Link trap	Enabled
Frame size	9198
Debounce time	0
Flow control	Disabled
Load interval	300
Ports, Link Flaps	
Admin mode	Disabled
Maximum count	5
Duration	10
Auto-Recovery Admin mode	Disabled
Auto-recovery interval	300
LAGs	
Auto-LAG Admin mode	Enabled
Auto-LAG Global Hash mode	2 Dest MAC, VLAN, EType, incoming port
Auto-LAG Global ABV Profile mode	Enabled

Feature Name/Setting	Default
Non Unicast Global Hash mode	Source MAC/IP address: Enabled Destination MAC/IP address: Enabled
Lag name	ch<n> where n is 1 to 64
Admin mode	Enabled
Hash mode	2 Dest MAC, VLAN, EType, incoming port
STP mode	Enabled
Static mode	Disabled
Link trap	Disabled
Local Preference Mode	Disabled
802.1AS, Global	
802.1AS status	Disabled
Local clock priority2	248
802.1AS, Interface	
Admin mode	Disabled
Pdelay threshold	3000
Allowed lost responses	3
Initial sync interval	–3
Initial Pdelay interval	0
Initial announce interval	0
SyncRx timeout	3
AnnounceRx timeout	3
MRP, Global	
MVRP mode	Disabled
MMRP mode	Disabled
MSRP mode	Disabled
MSRP talker pruning	Disabled
Periodic state machine (MVRP mode)	Enabled
Periodic state machine (MMRP mode)	Enabled

Feature Name/Setting	Default
MSRP maximum fan-in ports	0
MSRP boundary propagation	Disabled
MSRP PDU transmit time gap	100
802.1qav Mapping, Global	
EAV priority	Class A: 3 Class B: 2
EAV remap priority	Class A: 0 Class B: 0
MRP, Interface	
MVRP mode	Disabled
MMRP mode	Disabled
MSRP mode	Disabled
Join timer	20
Leave timer	100
Leave all timer	1000
MSRP SR class PVID	2
Qav Parameters, Interface	
Class A, MSRP delta bandwidth	100
Class B, MSRP delta bandwidth	0
L2 Loop Protection, Global	
Admin mode	Disabled
TLV advertised interval	5
Maximum PDU receive	1
L2 Loop Protection, Interface	
Keep alive	Disabled
RX action	Disabled

Routing, OSPF, OSPFv3, and multicast features default settings

The following table describes the default settings for the routing, OSPF, OSPFv3, and multicast features that you can configure.
(View-only settings are not included in the table but might be included in the data sheet, which you can download by visiting netgear.com/support/download/.)

Table 208.Routing, OSPF, OSPFv3, and multicast features default settings

Feature Name/Setting	Default
Routing Table, Route Preference	
Static	1
RIP	120
OSPF Intra	110
OSPF Inter	110
OSPF External	110
IP, Global	
Routing mode	Enabled
ICMP echo replies	64
ICMP redirects	Disabled
ICMP rate limit interval	1000
ICMP rate limit burst size	100
Global default gateway	Not configured
IP, Interface	
IP address configuration method	None
IP address	None
Subnet mask	None
Routing mode	Disabled
Administrative mode	Enabled
Forward net directed broadcasts	Disabled
Encapsulation type	Ethernet

Feature Name/Setting	Default
Proxy ARP	Enabled
Local proxy ARP	Disabled
Bandwidth	Depends on the port capacity (The bandwidth is automatically assigned but you can change it.)
ICMP destination unreachables	Enabled
ICMP redirects	Disabled
IP MTU	1500
IP, Secondary IP	
Secondary IP address	None
Secondary IP subnet mask	None
IPv6, Global	
IPv6 unicast routing	Disabled
Hop limit	64
ICMPv6 rate limit error interval	1000
ICMPv6 rate limit burst size	100
IPv6, Interface	
IPv6 mode	Disabled
DHCPv6 Client mode	Disabled
Stateless Address AutoConfig mode	Disabled
Routing mode	Disabled
Admin mode	Enabled
MTU	1500
Duplicate address detection transmits	1
Life time interval	1800
Adv NS interval	0
Adv reachable interval	0
Adv interval	600

Feature Name/Setting	Default
Adv managed config flag	Disabled
Adv other config flag	Disabled
Router preference	Medium
Adv suppress flag	Disabled
Destination unreachables	Enabled
IPv6, Route Preference	
Static	1
OSPFv3 intra	110
OSPFv3 inter	110
OSPFv3 external	110
ARP, Global	
Age time	1200
Response time	1
Retries	4
Cache size	4096
Dynamic renewal	Enabled
RIP, Global	
RIP admin mode	Enabled
Split Horizon mode	Simple
Auto Summary mode	Disabled
Host Routes Accept mode	Enabled
Default information originate	Disabled
Default metric	0
RIP, Interface	
Send version	RIP-2
Receive version	Both
RIP mode	Disabled
Authentication type	None

Feature Name/Setting	Default
Authentication Key ID	0
RIP, Route Redistribution	
Source	Connected
Redistribute mode	Disabled
Metric	0
Distribute list	0
OSPF, Default Route Advertise Configuration	
Default information originate	Disabled
Always	False
Metric	0
Metric type	External Type 2
OSPF, Global	
Admin mode	Enabled
Router ID	0.0.0.0
RFC 1583 compatibility	Enabled
Opaque LSA status	Enabled
Exit overflow interval (secs)	0
SPF delay time (secs)	5
SPF hold time (secs)	10
External LSDB limit	–1
Default metric	0
Maximum paths	16
AutoCost reference bandwidth	100
Default passive setting	Disabled
OSPF, Interface Configuration	
Area ID	0
Admin mode	Disabled
Router priority	1

Feature Name/Setting	Default
Retransmit interval	5
Hello interval	10
Dead interval	40
Iftransit delay interval	1
MTU ignore	Disabled
Passive mode	Disabled
Network type	Broadcast
Aythentication type	None
OSPF, Route Distribution	
Redistribute (for all sources)	Disabled
Metric (for all sources)	0
OSPF, NSF OSPF	
Support mode	Disabled
Restart interval	120
Helper support mode	Always
Helper strict LSA checking	Enabled
OSPFv3, Default Route Advertise Configuration	
Default information originate	Disabled
Always	False
Metric	0
Metric type	External Type 2
OSPFv3, Global	
Admin mode	Enabled
Router ID	0.0.0.0
Exit overflow interval (secs)	0
External LSDB limit	–1
Default metric	0
Maximum paths	16

Feature Name/Setting	Default
AutoCost reference bandwidth	100
Default passive setting	Disabled
Helper support mode	Always
Helper strict LSA checking	Enabled
OSPFv3, Interface Configuration	
Area ID	0
Admin mode	Disabled
Router priority	1
Retransmit interval	5
Hello interval	10
Dead interval	40
Iftransit delay interval	1
MTU ignore	Disabled
Passive mode	Disabled
Network type	Broadcast
Metric cost	1
OSPFv3, Route Distribution	
Redistribute (for all sources)	Disabled
Metric (for all sources)	0
Metric type (for all sources)	External Type 2
Tag (for all sources)	0
OSPFv3, NSF OSPF	
Support mode	Disabled
Restart interval	120
Router Discovery	
Advertise mode	Disabled
Advertise address	224.0.0.1
Maximum advertise interval	600

Feature Name/Setting	Default
Minimum advertise interval	450
Advertise lifetime	1800
Preference level	0
VRRP, Global	
Admin mode	Disabled
Multicast, Global	
Admin mode	Disabled
Multicast, Interface	
TTL threshold	1
Multicast, DVMRP, Global	
Admin mode	Disabled
Multicast, DVMRP, Interface	
Interface mode	Disabled
Interface metric	1
Multicast, IGMP, Global	
Admin mode	Disabled
Multicast, IGMP, Routing Interface	
Admin mode	Disabled
Version	V3
Robustness	2
Query interval	125
Query maximum response time	100
Startup query interval	31
Startup query count	2
Last member query interval	10
Last member query count	2
Multicast, IGMP, Proxy Interface	
Admin mode	Disabled

Feature Name/Setting	Default
Unsolicited report interval	1
Multicast, PIM, Global	
Admin mode	Disabled
Multicast, PIM, Interface	
Admin mode	Disabled
Hello interval	30
Join/Prune interval	60
BSR border	Disabled
DR priority	1
Multicast, PIM, BSR Candidate	
Hash mask length	30
Priority	0
Advertisement interval (secs)	60
Multicast IPv6, PIM, Global	
Admin mode	Disabled
Multicast IPv6, PIM, Interface	
Admin mode	Disabled
Hello interval	30
Join/Prune interval	60
BSR border	Disabled
DR priority	1
Multicast IPv6, PIM, BSR Candidate	
Hash mask length	126
Priority	0
Advertisement interval	60
Multicast IPv6, MLD, Global	
Admin mode	Disabled
Multicast IPv6, MLD, Routing Interface	

Feature Name/Setting	Default
Admin mode	Disabled
Version	V2
Query interval	125
Query maximum response time	10000
Startup query interval	31
Startup query count	2
Last member query interval	1000
Last member query count	2
Multicast IPv6, MLD, Proxy Interface	
Admin mode	Disabled
Unsolicited report interval	1

QoS features default settings

The following table describes the default settings for the QoS features that you can configure. (View-only settings are not included in the table but might be included in the data sheet, which you can download by visiting netgear.com/support/download/.)

Table 209.QoS features default settings

Feature Name/Setting	Default
CoS, Global	
Global trust mode	trust IP-DSCP applied to all interfaces
CoS, 802.1p Queue Mapping	
802.1p priority (802.1p -> queue)	0 -> 1 1 -> 0 2 -> 0 3 -> 1 4 -> 2 5 -> 2 6 -> 3 7 -> 3
CoS, IP DSCP Queue Mapping	
IP DSCP priority (IP DSCP -> queue)	0 through 7 -> 1

Feature Name/Setting	Default
	8 through 23 -> 0 24 through 31 -> 1 32 through 47 -> 2 48 through 63 -> 3
CoS, Interface	
Interface trust mode	IP DSCP
Interface shaping rate	0
CoS, Interface Queue	
Queue ID	0
Minimum bandwidth	0
Scheduler type	Weighted
Queue management type	Taildrop
CoS, Queue Drop Preference	
Drop precedence level	1 (for all interfaces and all queues)
WRED minimum threshold	40 (for all interfaces and all queues)
WRED maximum threshold	100 (for all interfaces and all queues)
WRED drop probability scale	10 (for all interfaces and all queues)
DiffServ, Global No default classes, policies, and services configured	
DiffServ admin mode	Enabled

Security features default settings

The following table describes the default settings for the security features that you can configure. (View-only settings are not included in the table but might be included in the data sheet, which you can download by visiting netgear.com/support/download/.)

Table 210.Security features default settings

Feature Name/Setting	Default
Management Security, Local User, User Password	
Password minimum length	8
Password aging (days)	0

Feature Name/Setting	Default
Password history	0
Lockout attempts	5
Unlock time	5
Unlock timer mode	Enabled
Management Security, Local User, MFA Global	
Multifactor authentication mode	Disabled
Management Security, Enable Password	
Encryption type	SHA512
Management Security, Line Password	
Encryption type (Console, Telnet, and SSH)	SHA512
Management Security, RADIUS, Global No RADIUS servers configured	
Source interface	VLAN 1
Maximum number of retransmits	4
Timeout duration	5
Accounting mode	Disabled
Radius Attribute 4 mode	Disabled
Management Security, TACAS+, Global No TACACS+ servers configured	
Connection timeout	5
Source interface	VLAN 1
Management Security, Authentication List	
defaultList	Option 1 is Local, options 2 through 6 are N/A
networkList	Option 1 is Local, options 2 through 6 are N/A
Management Security, Enable Authentication List	
enableList	Option 1 is Enabled, option 2 is None, options 3, 4, and 5 are N/A
enableNetList	Option 1 is Enabled, option 2 is None, options 3, 4, and 5 are N/A

Feature Name/Setting	Default
Management Security, Dot1x Authentication List	
dot1xList	Disabled
Management Security, HTTP Authentication List	
httpList	Disabled
Management Security, HTTPS Authentication List	
httpsList	Disabled
Access, HTTP	
HTTP access	Enabled
HTTP port	49151
HTTP session soft timeout (minutes)	15
HTTP session hard timeout (hours)	24
Maximum number of HTTP sessions	16
Access, HTTPS	
HTTPS access	Disabled
HTTPS port	443
HTTPS session soft timeout (minutes)	15
HTTPS session hard timeout (hours)	24
Maximum number of HTTPS sessions	16
Access, SSH	
SSH admin mode	Disabled
SSH session timeout	5
Maximum number of SSH sessions	5
Authentication list	networkList
Enable Authentication list	enableList
SSH port	22
SSH public key authentication mode	Disabled
SCP server admin mode	Disabled
Max SSH authentication retries	3

Feature Name/Setting	Default
Access, Telnet, Authentication List	
Authentication list	networkList
Enable authentication list	enableList
Access, Telnet, Inbound Telnet	
Telnet server admin mode	Enabled
Allow new Telnet sessions	Enabled
Session timeout (minutes)	5
Maximum number of sessions	5
Access, Telnet, Outbound Telnet	
Allow new Telnet sessions	Enabled
Session timeout (minutes)	5
Maximum number of sessions	5
Access, Console	
Serial port login timeout (minutes)	5
Baud rate (bps)	115200
Authentication list	defaultList
Enable authentication list	enableList
Access, Denial of Service	
Denial of Service minimum TCP header size	20
Denial of Service ICMPv4	Disabled
Denial of Service maximum ICMPv4 packet size	512 (if maximum ICMPv4 packet size is enabled)
Denial of Service ICMPv6	Disabled
Denial of Service maximum ICMPv6 packet size	512 (if maximum ICMPv6 packet size is enabled)
Denial of Service first fragment	Disabled
Denial of Service ICMP fragment	Disabled
Denial of Service SIP=DIP	Disabled
Denial of Service SMAC=DMAC	Disabled
Denial of Service TCP FIN & URG & PSH	Disabled

Feature Name/Setting	Default
Denial of Service TCP Flag & Sequence	Disabled
Denial of Service TCP fragment	Disabled
Denial of Service TCP offset	Disabled
Denial of Service TCP Port	Disabled
Denial of Service TCP SYN	Disabled
Denial of Service TCP SYN & FIN	Disabled
Denial of Service UDP port	Disabled
Stacked VLAN0 tag drop mode	Disabled
Access, Access Control, Access Profile No access profile configured	
Port Authentication, 802.1X, Global	
Dot1x admin mode	Disabled
Authentication admin mode	Disabled
VLAN assignment mode	Disabled
EAPOL flood mode	Disabled
Dynamic VLAN Creation mode	Disabled
Monitor mode	Disabled
Users	admin
Login	defaultList
Port Authentication, 802.1X, Interface	
Control mode	Auto
Host mode	Multi-Domain-Multi-Host
MAB	Disabled
MAB authentication type	N/A
Quiet period	60
Transmit period	30
Guest VLAN ID	0
Unauthenticated VLAN ID	0

Feature Name/Setting	Default
Maximum reauthentication requests	2
Supplicant timeout	30
Server timeout	30
Maximum requests	2
PAE capabilities	Authenticator
Periodic reauthentication	Disabled
Reauthentication period type	Server
Reauthentication period	0
Maximum users	48
Authentication order, Method 1	DOT1X
Authentication order, Method 2	MAB
Authentication order, Method 3	CAPTIVE-PORTAL
Traffic Control, MAC Filter No MAC filters configured	
Traffic Control, Port Security, Port Administration	
Port Security mode	Disabled
Traffic Control, Port Security, Interface	
Security mode	Disabled
Maximum learned MAC addresses	4096
Maximum static MAC address es	48
Violation shutdown	Disabled
Violation trap	Disabled
Traffic Control, Private Group	
Private groups	None configured
Traffic Control, Protected Port	
Protected ports	None configured
Traffic Control, Private VLAN	
Private VLANs	Unconfigured

Feature Name/Setting	Default
Traffic Control, Storm Control, Global	
Broadcast storm control all	Enabled
Multicast storm control all	Disabled
Unknown unicast storm control all	Disabled
Traffic Control, Storm Control, Interface	
Broadcast storm, Recovery mode	Enabled
Broadcast storm, Recovery level type	Percent
Broadcast storm, Recovery level	5
Broadcast storm, Control action	None
Multicast storm, Recovery mode	Disabled
Multicast storm, Recovery level type	Percent
Multicast storm, Recovery level	5
Multicast storm, Control action	None
Unicast storm, Recovery mode	Disabled
Unicast storm, Recovery level type	Percent
Unicast storm, Recovery level	5
Unicast storm, Control action	None
Control, DHCP Snooping, Global	
DHCP Snooping mode	Disabled
MAC address validation	Enabled
Control, DHCP Snooping, Interface	
Trust mode	Disabled
Invalid packets	Disabled
Rate limit (pps)	None
Burst interval (secs)	N/A
Control, DHCP Snooping, Persistent	
Store	Local
Write delay	300

Feature Name/Setting	Default
Control, DHCPv6 Snooping, Global	
DHCPv6 Snooping mode	Disabled
MAC address validation	Enabled
Control, DHCPv6 Snooping, Interface	
Trust mode	Disabled
Invalid packets	Disabled
Rate limit (pps)	None
Burst interval (secs)	N/A
Control, DHCPv6 Snooping, Persistent	
Store	Local
Write delay	300
Control, IP Source Guard, Interface	
IPSG mode	Disabled
IPSG port security	Disabled
Control, IPv6 Source Guard, Interface	
IPv6SG mode	Disabled
IPv6SG port security	Disabled
Control, Dynamic ARP Inspection, Global	
Validate source MAC	Disabled
Validate destination MAC	Disabled
Validate IP	Disabled
Control, Dynamic ARP Inspection, VLAN	
VLAN ID	1
Admin mode	Disabled
Invalid packets	Enabled
Static flag	Disabled
Control, Dynamic ARP Inspection, Interface	
Trust mode	Disabled

Feature Name/Setting	Default
Rate limit (pps)	15
Burst interval (secs)	1
Control, Captive Portal, Global	
Admin mode	Disabled
Additional HTTP port	0
Additional HTTP secure port	0
Authentication timeout	300
Control, Captive Portal, CP Configuration	
CP ID	1
CP Name	Default
Admin mode	Enabled
Protocol	HTTP
Verification	Guest
Block	Not Blocked
Group	0
Idle timeout	0
User logout	Disabled
RADIUS authentication server	None
Redirect mode	Disabled
Redirect URL	/v1/security/captive_portal/cp_welcome.html
Background color	#BFBFBF
Foreground color	#999999
Separator color	#B70024
Maximum bandwidth down	0
Maximum bandwidth down	0
Maximum input	0
Maximum output	0
Maximum total	0

Feature Name/Setting	Default
Control, Captive Portal, CP Group	
Group ID	1
Group name	Default
Control, Captive Portal, CP Trap Flags	
CP Trap mode	Disabled
Client authentication failure	None
Client connect	Disabled
Client DB full	None
Client disconnect	Disabled
ACL No default ACLs configured	
Total number of supported ACLs	100. This total number applies to all MAC ACLs, IPv4 ACLs, and IPv6 ACLs together.
MAC ACLs	None configured
IP ACLs	None configured
IPv6 ACLs	None configured

Monitoring features default settings

The following table describes the default settings for the monitoring features that you can configure.
(View-only settings are not included in the table but might be included in the data sheet, which you can download by visiting netgear.com/support/download/.)

Table 211.Monitoring features default settings

Feature	Default
Logs, Memory Log	
Admin status	Enabled
Behavior	Wrap
Severity filter	Notice
Threshold	80

Feature	Default
Logs, Command Log	
Admin status	Disabled
Logs, Console Log	
Admin status	Disabled
Severity filter	Error
Logs, Server Log	
Admin status	Disabled
Local UDP port	514
Source nterface	VLAN 1
Logs, USB Log	
USB log mode	Overwrite on Full
Severity filter	Notice
Mirroring, Multiple Port Mirroring, Global	
Session ID	1
Destination port	None
Admin mode	False
Filter type	None
Mirroring, Multiple Port Mirroring, Interface	
Direction	None
Mirroring, RSPAN VLAN	
VLAN ID	1
Admin mode	False
Mirroring, RSPAN Source Switch, Global	
Session ID	1
Admin mode	False
RSPAN destination VLAN	None
Filter type	None
Mirroring, RSPAN Source Switch, Interface	

Feature	Default
Direction	None
Mirroring, RSPAN Destination Switch, Global	
Session ID	1
Admin mode	False
RSPAN source VLAN	None
RSPAN destination Port	None
Filter type	None
sFlow Agent	
Source interface	VLAN 1
sFlow Receiver	
Receiver timeout	0
No timeout	False
Maximum datagram size	1400
Receiver address	0.0.0.0
Receiver port	6343
sFlow Interface	
Poller, Receiver index	0
Poller, Poller Interval	0
Sampler, Receiver index	0
Sampler, Sampling rate	0
Sampler, Maximum header size	128

General hardware technical specifications

The following hardware technical specifications apply to all M4350 series switch models:

Table 212.General hardware technical specifications for the M4350 series switch models

Feature	Description
Operating temperature	PoE+ and PoE++ models: 32° to 104°F (0° to 40°C)
	Non-PoE models: 32° to 113°F (0° to 45°C)

Feature	Description
Operating humidity	90% maximum relative humidity, noncondensing
Operating altitude	10,000 ft (3,000 m) maximum
Storage temperature	–4° to 158°F (–20° to 70°C)
Storage humidity	95% maximum relative humidity, noncondensing
Storage altitude	10,000 ft (3,000 m) maximum
Electromagnetic emissions and immunity certifications	Class A for FCC, ISED, CE, UKCA, VCCI, RCM, BSMI, KC
Safety certification	CB, CE, CSA, RCM, BSMI, KC (For PoE models only)

Model-specific hardware technical specifications

The tables in the following sections describes the unique main hardware technical specifications for each model.
For more hardware information, see the data sheet, which you can download by visiting netgear.com/support/download/.

M4350-8X8F (SKU XSM4316)

The following table shows the specifications that are specific to model M4350-8X8F.

Table 213.Hardware technical specifications for model M4350-8X8F

Feature		Description
Network interfaces		Eight 10G/5G/2.5G/1G/100M Ethernet ports Eight 1G/10G SFP+ fiber ports
AC power input		100–240V ~ 50–60Hz, 3A
Power consumption	Max. load Standby	81.1W 30.5W
Dimensions	Width Height Depth	8.66 in (220 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		8.93 lb (4.05 kg)

M4350-12X12F (SKU XSM4324)

The following table shows the specifications that are specific to model M4350-12X12F.

Table 214.Hardware technical specifications for model M4350-12X12F

Feature		Description
Network interfaces		Twelve 10G/5G/2.5G/1G/100M Ethernet ports Twelve 1G/10G SFP+ fiber ports
AC power input		100–240V ~ 50–60Hz, 3A
Power consumption	Max. load Standby	95.8W 34W
Dimensions	Width Height Depth	8.66 in (220 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		9.48 lb (4.3 kg)

M4350-24G4XF (SKU GSM4328)

The following table shows the specifications that are specific to model M4350-24G4XF.

Table 215.Hardware technical specifications for model M4350-24G4XF

Feature		Description
Network interfaces		Twenty-four PoE+ 1G/100M Ethernet ports Four 1G/10G SFP+ fiber ports
AC power input		100–240V ~ 50–60Hz, 10A
Power consumption	Max. PoE No PoE Standby	792.1W 80W 32W
PoE budget, internal power supply		648W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		14.13 lb (6.41 kg)

M4350-48G4XF (SKU GSM4352)

The following table shows the specifications that are specific to model M4350-48G4XF.

Table 216.Hardware technical specifications for model M4350-48G4XF

Feature		Description
Network interfaces		Forty-eight PoE+ 1G/100M Ethernet ports Four 1G/10G SFP+ fiber ports
AC power input		100–240V ~ 50–60Hz, 10A
Power consumption	Max. PoE No PoE Standby	384W 99W 48.5W
PoE budget, internal power supply		236W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		15.85 lb (7.19 kg)

M4350-24X4V (SKU XSM4328CV)

The following table shows the specifications that are specific to model M4350-24X4V.

Table 217.Hardware technical specifications for model M4350-24X4V

Feature		Description
Network interfaces		Twenty-four PoE+ 10G/5G/2.5G/1G/100M Ethernet ports Four 25G/10G SFP28 fiber ports
AC power input		100–240V ~ 50–60Hz, 10A
Power consumption	Max. PoE No PoE Standby	750.1W 119.4W 53.6W
PoE budget, internal power supply		576W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		14.51 lb (6.58 kg)

M4350-24F4V (SKU XSM4328FV)

The following table shows the specifications that are specific to model M4350-24F4V.

Table 218.Hardware technical specifications for model M4350-24F4V

Feature		Description
Network interfaces		Twenty-four 1G/10G SFP+ fiber ports Four 25G/10G SFP28 fiber ports
AC power input		100–240V ~ 50–60Hz, 3A
Power consumption	Max. load Standby	119.3W 45.2W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		13.78 lb (6.25 kg)

M4350-44M4X4V (SKU MSM4352)

The following table shows the specifications that are specific to model M4350-44M4X4V.

Table 219.Hardware technical specifications for model M4350-44M4X4V

Feature		Description
Network interfaces		Forty-four PoE++ 2.5G/1G/100M Ethernet ports Four PoE++ 10G/5G/2.5G/1G/100M Ethernet ports Four 25G/10G SFP28 fiber ports
AC power input		100–240V ~ 50–60Hz, 10A
Power consumption	Max. PoE No PoE Standby	351.5W 133.5W 56.5W
PoE budget, internal power supply		194W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		16.18 lb (7.34 kg)

M4350-36X4V (SKU XSM4340CV)

The following table shows the specifications that are specific to model M4350-36X4V.

Table 220.Hardware technical specifications for model M4350-36X4V

Feature		Description
Network interfaces		Thirty-six PoE++ 10G/5G/2.5G/1G/100M Ethernet ports Four 25G/10G SFP28 fiber ports
AC power input		100–240V ~ 50–60Hz, 12A
Power consumption	Max. PoE No PoE Standby	454W 148W 59.3W
PoE budget, internal power supply		280W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		16.16 lb (7.33 kg)

M4350-24X8F8V (SKU XSM4340V)

The following table shows the specifications that are specific to model M4350-24X8F8V.

Table 221.Hardware technical specifications for model M4350-24X8F8V

Feature		Description
Network interfaces		Twenty-four PoE++ 10G/5/2.5G/1G/100M Ethernet ports Eight 10G/1G SFP+ fiber ports Eight 25G/10G SFP28 fiber ports
AC power input		100–240V ~ 50–60Hz, 12A
Power consumption	Max. PoE No PoE Standby	464W 151.8W 53.4W
PoE budget, internal power supply		290W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		15.96 lb (7.24 kg)

M4350-32F8V (SKU XSM4340FV)

The following table shows the specifications that are specific to model M4350-32F8V.

Table 222.Hardware technical specifications for model M4350-32F8V

Feature		Description
Network interfaces		Thirty-two 10G/1G SFP+ fiber ports Eight 25G/10G SFP28 fiber ports
AC power input		100–240V ~ 50–60Hz, 12A
Power consumption	Max. load Standby	156.4W 42.6W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		15.32 lb (6.95 kg)

M4350-16V4C (SKU VSM4320C)

The following table shows the specifications that are specific to model M4350-16V4C.

Table 223.Hardware technical specifications for model M4350-16V4C

Feature		Description
Network interfaces		Sixteen 25G/10G SFP28 fiber ports Four 40G/100G QSFP28 fiber ports
AC power input		100–240V ~ 50–60Hz, 12A
Power consumption	Max. load Standby	143.4W 60.4W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		15.77 lb (7.15 kg)

M4350-40X4C (SKU XSM4344C)

The following table shows the specifications that are specific to model M4350-40X4C.

Table 224.Hardware technical specifications for model M4350-40X4C

Feature	Description
Network interfaces	Fourty PoE++ 10G/5G/2.5G/1G/100M Ethernet ports Four 40G/100G QSFP28 fiber ports

Feature		Description
AC power input		100–240V ~ 50–60Hz, 12A
Power consumption	Max. PoE No PoE Standby	413W 200.4W 98.8W
PoE budget, internal power supply		196W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		17.11 lb (7.76 kg)

M4350-8M2V (SKU MSM4310)

The following table shows the specifications that are specific to model M4350-8M2V.

Table 225.Hardware technical specifications for model M4350-8M2V

Feature		Description
Network interfaces		Eight PoE++ 2.5G/1G/100M Ethernet ports Two 25G/10G SFP28 fiber ports
AC power input		100–240V ~ 50–60Hz, 3A
Power consumption	Max. load Standby	82W 61.2W
Dimensions	Width Height Depth	8.66 in (220 mm) 1.70 in (43.2 mm), 1U 11.81 in (300 mm)
Weight		6.39 lb (2.9 kg)

M4350-24M4X4V (SKU MSM4332)

The following table shows the specifications that are specific to model M4350-24M4X4V.

Table 226.Hardware technical specifications for model M4350-24M4X4V

Feature	Description
Network interfaces	Twenty-four PoE++ 2.5G/1G/100M Ethernet ports Four PoE++ 10G/5G/2.5G/1G/100M Ethernet ports Four 25G/10G SFP28 fiber ports

Feature		Description
AC power input		100–240V ~ 50–60Hz, 7.5-3A
Power consumption	Max. load Standby	123.4W 53W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		16.26 lb (7.38 kg)

M4350-24F4X (SKU MSM4328F)

The following table shows the specifications that are specific to model M4350-24F4X.

Table 227.Hardware technical specifications for model M4350-24F4X

Feature		Description
Network interfaces		Twenty-four 2.5G/1G/100M Ethernet ports Four combo multispeed 10G/5G/2.5G/1G/100M Ethernet ports Four combo SFP+ 10G/1G fiber ports
AC power input		100–240V ~ 50–60 Hz, 6.3-3A
Power consumption	Max. load Standby	115.5W 39W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		16.26 lb (7.38 kg)

M4350-40F4C (SKU XSM4344FC)

The following table shows the specifications that are specific to model M4350-40F4C.

Table 228.Hardware technical specifications for model M4350-40F4C

Feature	Description
Network interfaces	Forty 1G/10G SFP+ fiber ports Four 40G/100G QSFP28 fiber ports
AC power input	100–240V ~ 50–60Hz, 6.3-3A

Feature		Description
Power consumption	Max. load Standby	213.6W 65W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		16.57 lb (7.52 kg)

M4350-16M4V (SKU MSM4320)

The following table shows the specifications that are specific to model M4350-16M4V.

Table 229.Hardware technical specifications for model M4350-16M4V

Feature		Description
Network interfaces		Eight PoE++ 2.5G/1G/100M multispeed ports with Neutrik connectors Eight PoE++ 2.5G/1G/100M multispeed Ethernet RJ-45 ports
AC power input		100–240V ~ 50–60Hz, 8-3A
Power consumption	Max. load Standby	109.6W 50.6W
Dimensions	Width Height Depth	17.32 in (440 mm) 1.70 in (43.2 mm), 1U 17.01 in (432 mm)
Weight		17.31 lb (7.85 kg)

M4350-16C (SKU CSM4316)

The following table shows the specifications that are specific to model M4350-16C.

Table 230.Hardware technical specifications for model M4350-16C

Feature		Description
Network interfaces		16 100G/40G QSFP28 fiber uplink ports
AC power input		100–240V ~ 50–60Hz, 2A
Power consumption	Max. load Standby	161.9W 69.1W
Dimensions	Width	17.32 in (440 mm)

Feature		Description
	Height Depth	1.70 in (43.2 mm), 1U 15.75 in (400 mm)
Weight		17.70 lb (8.03 kg)