



10G/40G Top-of-Rack Switches

AS5700-54X

AS6700-32X

| Software Release v1.1.163.153

Web Management Guide

Web Management Guide

AS5700-54X

54-Port 10G Data Center Switch
with 48 10GBASE SFP+ Ports,
6 40GBASE QSFP Ports,
2 Power Supply Units,
and 5 Fan Trays (5 Fans – F2B and B2F Airflow)

AS6700-32X

32-Port 40G Data Center Switch
with 20 40G QSFP+ Ports,
2 40G Expansion Slots,
2 Power Supply Units,
and 5 Fan Trays (5 Fans – F2B or B2F Airflow)

How to Use This Guide

This guide includes detailed information on the switch software, including how to operate and use the management functions of the switch. To deploy this switch effectively and ensure trouble-free operation, you should first read the relevant sections in this guide so that you are familiar with all of its software features.

Who Should Read This Guide? This guide is for network administrators who are responsible for operating and maintaining network equipment. The guide assumes a basic working knowledge of LANs (Local Area Networks), the Internet Protocol (IP), and Simple Network Management Protocol (SNMP).

How This Guide is Organized This guide describes the switch's web browser interface. For more detailed information on the switch's key features refer to the *Administrator's Guide*.

The guide includes these sections:

- ◆ Section I **"Getting Started"** — Includes an introduction to switch management, and the basic settings required to access the management interface.
- ◆ Section II **"Web Configuration"** — Includes all management options available through the web browser interface.
- ◆ Section III **"Appendices"** — Includes information on troubleshooting switch management access.

Related Documentation This guide focuses on switch software configuration through the web browser.

For information on how to manage the switch through the command line interface, see the following guide:

CLI Reference Guide



Note: For a description of how to initialize the switch for management access via the CLI, web interface or SNMP, refer to "Initial Switch Configuration" in the *CLI Reference Guide*.

For information on how to install the switch, see the following guide:

Installation Guide

For all safety information and regulatory statements, see the following documents:

Quick Start Guide

Safety and Regulatory Information

Conventions The following conventions are used throughout this guide to show information:



Note: Emphasizes important information or calls your attention to related features or instructions.



Caution: Alerts you to a potential hazard that could cause loss of data, or damage the system or equipment.



Warning: Alerts you to a potential hazard that could cause personal injury.

Revision History This section summarizes the changes in each revision of this guide.

December 2015 Revision

This is the second version of this guide. This guide is valid for software release v1.1.163.153. It contains the following changes:

Table 1: Revision History

Description of Changes

Added:

- "IEEE 802.1Q Tunneling" on page 154
- "Issuing MAC Address Traps" on page 170
- "Rate Limiting" on page 195
- "Web Authentication" on page 243
- "Network Access (MAC Address Authentication)" on page 246
- "Configuring an ARP ACL" on page 279
- "Configuring ACL Mirroring" on page 282
- "Showing ACL Hardware Counters" on page 284
- "ARP Inspection" on page 285
- "Configuring Port Security" on page 295
- "Configuring 802.1X Port Authentication" on page 297
- "IPv4 Source Guard" on page 306
- "IPv6 Source Guard" on page 312

Table 1: Revision History (Continued)

Description of Changes
Added (Continued)
"Displaying Transceiver Data" on page 123
"DHCP Snooping" on page 318
"Connectivity Fault Management" on page 390
"MLD Snooping (Snooping and Query for IPv6)" on page 462
"Layer 3 IGMP (Query used with Multicast Routing)" on page 470
"Domain Name Service" on page 505
"Configuring the Routing Information Protocol" on page 544
"Specifying Passive Interfaces" on page 596
"Multicast Routing" on page 599
Updated:
Table 5, Switch Main Menu," on page 52
"Configuring a Dynamic Trunk" on page 130
"Configuring VLAN Groups" on page 147
"Creating CVLAN to SPVLAN Mapping Entries" on page 159
"Setting Static Addresses" on page 165
"Displaying Interface Settings for STA" on page 186
"Creating QoS Policies" on page 224
"Configuring Global Settings for HTTPS" on page 255
"Showing the Host Key Pair" on page 263
"Access Control Lists" on page 266
"Showing TCAM Utilization" on page 268
"Setting LLDP Timing Attributes" on page 332
"Configuring Local SNMPv3 Users" on page 367
"Setting IGMP Snooping Status per Interface" on page 445
"Setting the Switch's IP Address (IP Version 4)" on page 481
"Configuring the IPv6 Default Gateway" on page 485
Table 36, Show IPv6 Neighbors - display description," on page 495
"Dynamic Host Configuration Protocol" on page 511
"Using the Ping Function" on page 519
Deleted:

November 2015 Revision

This is the first version of this guide. This guide is valid for software release v1.1.0.152.

Contents

How to Use This Guide	3
Contents	7
Figures	19
Tables	33

Section I Getting Started 35

1 Introduction	37
Key Features	37
Description of Software Features	38
Equal-cost Multipath Load Balancing	42
System Defaults	44

Section II Web Configuration 47

2 Using the Web Interface	49
Connecting to the Web Interface	49
Navigating the Web Browser Interface	50
Home Page	50
Configuration Options	51
Panel Display	51
Main Menu	52
3 Basic Management Tasks	71
Displaying System Information	72
Displaying Hardware/Software Versions	73
Configuring Support for Jumbo Frames	74
Displaying Bridge Extension Capabilities	75

Managing System Files	77
Copying Files via FTP/TFTP or HTTP	77
Saving the Running Configuration to a Local File	79
Setting The Start-Up File	80
Showing System Files	80
Automatic Operation Code Upgrade	81
Setting the System Clock	85
Setting the Time Manually	85
Setting the SNTP Polling Interval	86
Configuring NTP	87
Configuring Time Servers	88
Setting the Time Zone	92
Configuring The Console Port	93
Configuring Telnet Settings	95
Displaying CPU Utilization	97
Displaying Memory Utilization	98
Resetting the System	98
4 Interface Configuration	103
Port Configuration	104
Configuring by Port List	104
Configuring by Port Range	106
Displaying Connection Status	107
Configuring Local Port Mirroring	108
Configuring Remote Port Mirroring	110
Showing Port or Trunk Statistics	114
Displaying Statistical History	118
Displaying Transceiver Data	122
Configuring Transceiver Thresholds	123
Trunk Configuration	125
Configuring a Static Trunk	126
Configuring a Dynamic Trunk	129
Displaying LACP Port Counters	135
Displaying LACP Settings and Status for the Local Side	136
Displaying LACP Settings and Status for the Remote Side	138

Configuring Load Balancing	139
Traffic Segmentation	141
Enabling Traffic Segmentation	141
Configuring Uplink and Downlink Ports	142
5 VLAN Configuration	145
IEEE 802.1Q VLANs	145
Configuring VLAN Groups	147
Adding Static Members to VLANs	150
IEEE 802.1Q Tunneling	154
Enabling QinQ Tunneling on the Switch	158
Creating CVLAN to SPVLAN Mapping Entries	159
Adding an Interface to a QinQ Tunnel	161
6 Address Table Settings	163
Configuring MAC Address Learning	163
Setting Static Addresses	165
Changing the Aging Time	167
Displaying the Dynamic Address Table	168
Clearing the Dynamic Address Table	169
Issuing MAC Address Traps	170
7 Spanning Tree Algorithm	173
Overview	173
Configuring Global Settings for STA	175
Displaying Global Settings for STA	180
Configuring Interface Settings for STA	181
Displaying Interface Settings for STA	186
Configuring Multiple Spanning Trees	188
Configuring Interface Settings for MSTP	192
8 Congestion Control	195
Rate Limiting	195
Storm Control	196

9	Class of Service	199
	Layer 2 Queue Settings	199
	Setting the Default Priority for Interfaces	199
	Selecting the Queue Mode	200
	Mapping CoS Values to Egress Queues	203
	Layer 3/4 Priority Settings	206
	Setting Priority Processing to IP Precedence/DSCP or CoS	206
	Mapping Ingress DSCP Values to Internal DSCP Values	207
	Mapping CoS Priorities to Internal DSCP Values	210
	Mapping Internal DSCP Values to Egress CoS Values	212
	Mapping IP Precedence Values to Internal DSCP Values	214
	Mapping IP Port Priority to Internal DSCP Values	216
10	Quality of Service	219
	Overview	219
	Configuring a Class Map	220
	Creating QoS Policies	224
	Attaching a Policy Map to a Port	233
11	Security Measures	235
	AAA Authentication, Authorization and Accounting	236
	Configuring Local/Remote Logon Authentication	237
	Configuring Remote Logon Authentication Servers	238
	Configuring User Accounts	241
	Web Authentication	243
	Configuring Global Settings for Web Authentication	244
	Configuring Interface Settings for Web Authentication	245
	Network Access (MAC Address Authentication)	246
	Configuring Global Settings for Network Access	248
	Configuring Network Access for Ports	249
	Configuring Port Link Detection	251
	Configuring a MAC Address Filter	252
	Displaying Secure MAC Address Information	254
	Configuring HTTPS	255
	Configuring Global Settings for HTTPS	255
	Replacing the Default Secure-site Certificate	257

Configuring the Secure Shell	258
Configuring the SSH Server	261
Generating the Host Key Pair	262
Showing the Host Key Pair	264
Importing User Public Keys	265
Access Control Lists	267
Showing TCAM Utilization	269
Setting the ACL Name and Type	270
Configuring a Standard IPv4 ACL	272
Configuring an Extended IPv4 ACL	273
Configuring a Standard IPv6 ACL	275
Configuring an Extended IPv6 ACL	277
Configuring a MAC ACL	278
Configuring an ARP ACL	280
Binding a Port to an Access Control List	282
Configuring ACL Mirroring	283
Showing ACL Hardware Counters	285
ARP Inspection	286
Configuring Global Settings for ARP Inspection	287
Configuring VLAN Settings for ARP Inspection	289
Configuring Interface Settings for ARP Inspection	290
Displaying ARP Inspection Statistics	292
Displaying the ARP Inspection Log	293
Filtering IP Addresses for Management Access	294
Configuring Port Security	296
Configuring 802.1X Port Authentication	298
Configuring 802.1X Global Settings	300
Configuring Port Authenticator Settings for 802.1X	301
Displaying 802.1X Statistics	305
IPv4 Source Guard	307
Configuring Ports for IPv4 Source Guard	307
Configuring Static Bindings for IP Source Guard	309
Displaying Information for Dynamic IPv4 Source Guard Bindings	312
IPv6 Source Guard	313
Configuring Ports for IPv6 Source Guard	313

Configuring Static Bindings for IPv6 Source Guard	316
Displaying Information for Dynamic IPv6 Source Guard Bindings	318
DHCP Snooping	319
DHCP Snooping Global Configuration	321
DHCP Snooping VLAN Configuration	323
Configuring Interfaces for DHCP Snooping	324
Displaying DHCP Snooping Binding Information	325
12 Basic Administration Protocols	327
Configuring Event Logging	327
System Log Configuration	327
Remote Log Configuration	330
Link Layer Discovery Protocol	331
Setting LLDP Timing Attributes	332
Configuring LLDP Interface Attributes	334
Configuring LLDP Interface Civic-Address	337
Displaying LLDP Local Device Information	339
Displaying LLDP Remote Device Information	343
Displaying Device Statistics	351
Simple Network Management Protocol	353
Configuring Global Settings for SNMP	355
Setting the Local Engine ID	356
Specifying a Remote Engine ID	357
Setting SNMPv3 Views	358
Configuring SNMPv3 Groups	361
Setting Community Access Strings	366
Configuring Local SNMPv3 Users	367
Configuring Remote SNMPv3 Users	369
Specifying Trap Managers	372
Creating SNMP Notification Logs	376
Showing SNMP Statistics	378
Remote Monitoring	380
Configuring RMON Alarms	380
Configuring RMON Events	383
Configuring RMON History Samples	385

Configuring RMON Statistical Samples	388
Connectivity Fault Management	390
Configuring Global Settings for CFM	394
Configuring Interfaces for CFM	397
Configuring CFM Maintenance Domains	398
Configuring CFM Maintenance Associations	403
Configuring Maintenance End Points	407
Configuring Remote Maintenance End Points	409
Transmitting Link Trace Messages	411
Transmitting Loop Back Messages	412
Transmitting Delay-Measure Requests	414
Displaying Local MEPs	416
Displaying Details for Local MEPs	417
Displaying Local MIPs	419
Displaying Remote MEPs	420
Displaying Details for Remote MEPs	421
Displaying the Link Trace Cache	423
Displaying Fault Notification Settings	425
Displaying Continuity Check Errors	425
UDLD Configuration	427
Configuring UDLD Protocol Intervals	427
Configuring UDLD Interface Settings	429
Displaying UDLD Neighbor Information	430
13 Multicast Filtering	433
Overview	433
IGMP Protocol	434
Layer 2 IGMP (Snooping and Query for IPv4)	435
Configuring IGMP Snooping and Query Parameters	437
Specifying Static Interfaces for an IPv4 Multicast Router	441
Assigning Interfaces to IPv4 Multicast Services	443
Setting IGMP Snooping Status per Interface	445
Filtering IGMP Query Packets	451
Displaying Multicast Groups Discovered by IGMP Snooping	452
Displaying IGMP Snooping Statistics	453

Filtering and Throttling IGMP Groups	457
Enabling IGMP Filtering and Throttling	457
Configuring IGMP Filter Profiles	458
Configuring IGMP Filtering and Throttling for Interfaces	460
MLD Snooping (Snooping and Query for IPv6)	462
Configuring MLD Snooping and Query Parameters	462
Setting Immediate Leave Status for MLD Snooping per Interface	464
Specifying Static Interfaces for an IPv6 Multicast Router	465
Assigning Interfaces to IPv6 Multicast Services	467
Showing MLD Snooping Groups and Source List	469
Layer 3 IGMP (Query used with Multicast Routing)	470
Configuring IGMP Proxy Routing	471
Configuring IGMP Interface Parameters	474
Configuring Static IGMP Group Membership	476
Displaying Multicast Group Information	478
14 IP Configuration	481
Setting the Switch's IP Address (IP Version 4)	481
Setting the Switch's IP Address (IP Version 6)	485
Configuring the IPv6 Default Gateway	485
Configuring IPv6 Interface Settings	486
Configuring an IPv6 Address	491
Showing IPv6 Addresses	494
Showing the IPv6 Neighbor Cache	495
Showing IPv6 Statistics	496
Showing the MTU for Responding Destinations	503
15 IP Services	505
Domain Name Service	505
Configuring General DNS Service Parameters	505
Configuring a List of Domain Names	506
Configuring a List of Name Servers	508
Configuring Static DNS Host to Address Entries	509
Displaying the DNS Cache	510
Dynamic Host Configuration Protocol	511
Specifying A DHCP Client Identifier	511

Configuring DHCP Relay Service	513
16 General IP Routing	515
Overview	515
Initial Configuration	515
IP Routing and Switching	516
Routing Path Management	517
Routing Protocols	518
Configuring IP Routing Interfaces	518
Configuring Local and Remote Interfaces	518
Using the Ping Function	519
Using the Trace Route Function	520
Address Resolution Protocol	522
ARP Timeout Configuration	522
Configuring Static ARP Addresses	523
Displaying Dynamic or Local ARP Entries	525
Displaying ARP Statistics	525
Configuring Static Routes	526
Displaying the Routing Table	528
Equal-cost Multipath Routing	529
17 Configuring Router Redundancy	533
Configuring VRRP Groups	534
Displaying VRRP Global Statistics	540
Displaying VRRP Group Statistics	541
18 Unicast Routing	543
Overview	543
Configuring the Routing Information Protocol	544
Configuring General Protocol Settings	545
Clearing Entries from the Routing Table	548
Specifying Network Interfaces	549
Specifying Passive Interfaces	551
Specifying Static Neighbors	552
Configuring Route Redistribution	553
Specifying an Administrative Distance	555

Configuring Network Interfaces for RIP	556
Displaying RIP Interface Settings	560
Displaying Peer Router Information	561
Resetting RIP Statistics	561
Configuring the Open Shortest Path First Protocol (Version 2)	562
Defining Network Areas Based on Addresses	563
Configuring General Protocol Settings	566
Displaying Administrative Settings and Statistics	569
Adding an NSSA or Stub	571
Configuring NSSA Settings	572
Configuring Stub Settings	575
Displaying Information on NSSA and Stub Areas	577
Configuring Area Ranges (Route Summarization for ABRs)	578
Redistributing External Routes	580
Configuring Summary Addresses (for External AS Routes)	582
Configuring OSPF Interfaces	584
Configuring Virtual Links	589
Displaying Link State Database Information	592
Displaying Information on Neighboring Routers	595
Specifying Passive Interfaces	596
19 Multicast Routing	599
Overview	599
Configuring Global Settings for Multicast Routing	602
Enabling Multicast Routing Globally	602
Displaying the Multicast Routing Table	603
Configuring PIM for IPv4	607
Enabling PIM Globally	607
Configuring PIM Interface Settings	608
Displaying PIM Neighbor Information	613
Configuring Global PIM-SM Settings	614
Configuring a PIM BSR Candidate	615
Configuring a PIM Static Rendezvous Point	617
Configuring a PIM RP Candidate	618
Displaying the PIM BSR Router	620

Displaying PIM RP Mapping	622
Configuring PIMv6 for IPv6	623
Enabling PIMv6 Globally	623
Configuring PIMv6 Interface Settings	624
Displaying PIM6 Neighbor Information	629
Configuring Global PIM6-SM Settings	630
Configuring a PIM6 BSR Candidate	631
Configuring a PIM6 Static Rendezvous Point	633
Configuring a PIM6 RP Candidate	635
Displaying the PIM6 BSR Router	637
Displaying RP Mapping	638

Section III	Appendices	641
	A Software Specifications	643
	Software Features	643
	Management Features	644
	Standards	645
	Management Information Bases	646
	B Troubleshooting	649
	Problems Accessing the Management Interface	649
	Using System Logs	650
	C License Information	651
	The GNU General Public License	651
	Glossary	655
	Index	663

Figures

Figure 1: Home Page	50
Figure 2: Front Panel Indicators	51
Figure 3: System Information	72
Figure 4: General Switch Information	73
Figure 5: Configuring Support for Jumbo Frames	75
Figure 6: Displaying Bridge Extension Configuration	76
Figure 7: Copy Firmware	78
Figure 8: Saving the Running Configuration	79
Figure 9: Setting Start-Up Files	80
Figure 10: Displaying System Files	81
Figure 11: Configuring Automatic Code Upgrade	84
Figure 12: Manually Setting the System Clock	86
Figure 13: Setting the Polling Interval for SNTP	87
Figure 14: Configuring NTP	88
Figure 15: Specifying SNTP Time Servers	89
Figure 16: Adding an NTP Time Server	90
Figure 17: Showing the NTP Time Server List	90
Figure 18: Adding an NTP Authentication Key	91
Figure 19: Showing the NTP Authentication Key List	91
Figure 20: Setting the Time Zone	93
Figure 21: Console Port Settings	95
Figure 22: Telnet Connection Settings	96
Figure 23: Displaying CPU Utilization	97
Figure 24: Displaying Memory Utilization	98
Figure 25: Restarting the Switch (Immediately)	100
Figure 26: Restarting the Switch (In)	101
Figure 27: Restarting the Switch (At)	101
Figure 28: Restarting the Switch (Regularly)	101
Figure 29: Configuring Connections by Port List	106

Figure 30: Configuring Connections by Port Range	107
Figure 31: Displaying Port Information	108
Figure 32: Configuring Local Port Mirroring	108
Figure 33: Configuring Local Port Mirroring	109
Figure 34: Displaying Local Port Mirror Sessions	109
Figure 35: Configuring Remote Port Mirroring	110
Figure 36: Configuring Remote Port Mirroring (Source)	113
Figure 37: Configuring Remote Port Mirroring (Intermediate)	113
Figure 38: Configuring Remote Port Mirroring (Destination)	114
Figure 39: Showing Port Statistics (Table)	117
Figure 40: Showing Port Statistics (Chart)	118
Figure 41: Configuring a History Sample	120
Figure 42: Showing Entries for History Sampling	120
Figure 43: Showing Status of Statistical History Sample	121
Figure 44: Showing Current Statistics for a History Sample	121
Figure 45: Showing Ingress Statistics for a History Sample	122
Figure 46: Displaying Transceiver Data	123
Figure 47: Configuring Transceiver Thresholds	125
Figure 48: Configuring Static Trunks	126
Figure 49: Creating Static Trunks	127
Figure 50: Adding Static Trunks Members	128
Figure 51: Configuring Connection Parameters for a Static Trunk	128
Figure 52: Showing Information for Static Trunks	129
Figure 53: Configuring Dynamic Trunks	129
Figure 54: Configuring the LACP Aggregator Admin Key	132
Figure 55: Enabling LACP on a Port	133
Figure 56: Configuring LACP Parameters on a Port	133
Figure 57: Showing Members of a Dynamic Trunk	134
Figure 58: Configuring Connection Settings for Dynamic Trunks	134
Figure 59: Displaying Connection Parameters for Dynamic Trunks	135
Figure 60: Displaying LACP Port Counters	136
Figure 61: Displaying LACP Port Internal Information	137
Figure 62: Displaying LACP Port Remote Information	139
Figure 63: Configuring Load Balancing	140
Figure 64: Enabling Traffic Segmentation	142

Figure 65: Configuring Members for Traffic Segmentation	143
Figure 66: Showing Traffic Segmentation Members	144
Figure 67: VLAN Compliant and VLAN Non-compliant Devices	146
Figure 68: Creating Static VLANs	148
Figure 69: Modifying Settings for Static VLANs	149
Figure 70: Showing Static VLANs	149
Figure 71: Configuring Static Members by VLAN Index	152
Figure 72: Configuring Static VLAN Members by Interface	153
Figure 73: Configuring Static VLAN Members by Interface Range	153
Figure 74: QinQ Operational Concept	155
Figure 75: Enabling QinQ Tunneling	159
Figure 76: Configuring CVLAN to SPVLAN Mapping Entries	160
Figure 77: Showing CVLAN to SPVLAN Mapping Entries	160
Figure 78: Adding an Interface to a QinQ Tunnel	162
Figure 79: Configuring MAC Address Learning	164
Figure 80: Configuring Static MAC Addresses	166
Figure 81: Displaying Static MAC Addresses	166
Figure 82: Setting the Address Aging Time	167
Figure 83: Displaying the Dynamic MAC Address Table	169
Figure 84: Clearing Entries in the Dynamic MAC Address Table	170
Figure 85: Issuing MAC Address Traps (Global Configuration)	171
Figure 86: Issuing MAC Address Traps (Interface Configuration)	171
Figure 87: STP Root Ports and Designated Ports	174
Figure 88: MSTP Region, Internal Spanning Tree, Multiple Spanning Tree	174
Figure 89: Spanning Tree – Common Internal, Common, Internal	175
Figure 90: Configuring Global Settings for STA (STP)	179
Figure 91: Configuring Global Settings for STA (RSTP)	179
Figure 92: Configuring Global Settings for STA (MSTP)	180
Figure 93: Displaying Global Settings for STA	181
Figure 94: Determining the Root Port	183
Figure 95: Configuring Interface Settings for STA	186
Figure 96: STA Port Roles	187
Figure 97: Displaying Interface Settings for STA	188
Figure 98: Creating an MST Instance	190
Figure 99: Displaying MST Instances	190

Figure 100: Modifying the Priority for an MST Instance	191
Figure 101: Displaying Global Settings for an MST Instance	191
Figure 102: Adding a VLAN to an MST Instance	192
Figure 103: Displaying Members of an MST Instance	192
Figure 104: Configuring MSTP Interface Settings	194
Figure 105: Displaying MSTP Interface Settings	194
Figure 106: Configuring Rate Limits	196
Figure 107: Configuring Storm Control	197
Figure 108: Setting the Default Port Priority	200
Figure 109: Setting the Queue Mode (Strict)	202
Figure 110: Setting the Queue Mode (WRR)	202
Figure 111: Setting the Queue Mode (Strict and WRR)	203
Figure 112: Mapping CoS Values to Egress Queues	205
Figure 113: Showing CoS Values to Egress Queue Mapping	205
Figure 114: Setting the Trust Mode	207
Figure 115: Configuring DSCP to DSCP Internal Mapping	209
Figure 116: Showing DSCP to DSCP Internal Mapping	209
Figure 117: Configuring CoS to DSCP Internal Mapping	211
Figure 118: Showing CoS to DSCP Internal Mapping	212
Figure 119: Configuring DSCP to CoS Egress Mapping	213
Figure 120: Showing DSCP to CoS Egress Mapping	214
Figure 121: Configuring IP Precedence to DSCP Internal Mapping	215
Figure 122: Showing the IP Precedence to DSCP Internal Map	216
Figure 123: Configuring IP Port Number to DSCP Internal Mapping	217
Figure 124: Showing IP Port Number to DSCP Internal Mapping	218
Figure 125: Configuring a Class Map	221
Figure 126: Showing Class Maps	222
Figure 127: Adding Rules to a Class Map	223
Figure 128: Showing the Rules for a Class Map	223
Figure 129: Configuring a Policy Map	231
Figure 130: Showing Policy Maps	231
Figure 131: Adding Rules to a Policy Map	232
Figure 132: Showing the Rules for a Policy Map	233
Figure 133: Attaching a Policy Map to a Port	234
Figure 134: Configuring the Authentication Sequence	238

Figure 135: Authentication Server Operation	238
Figure 136: Configuring Remote Authentication Server (RADIUS)	241
Figure 137: Configuring Remote Authentication Server (TACACS+)	241
Figure 138: Configuring User Accounts	243
Figure 139: Showing User Accounts	243
Figure 140: Configuring Global Settings for Web Authentication	245
Figure 141: Configuring Interface Settings for Web Authentication	246
Figure 142: Configuring Global Settings for Network Access	249
Figure 143: Configuring Interface Settings for Network Access	251
Figure 144: Configuring Link Detection for Network Access	252
Figure 145: Configuring a MAC Address Filter for Network Access	253
Figure 146: Showing the MAC Address Filter Table for Network Access	254
Figure 147: Showing Addresses Authenticated for Network Access	255
Figure 148: Configuring HTTPS	257
Figure 149: Downloading the Secure-Site Certificate	258
Figure 150: Configuring the SSH Server	262
Figure 151: Generating the SSH Host Key Pair	263
Figure 152: Showing the SSH Host Key Pair	264
Figure 153: Showing the SSH Host Key Pair	265
Figure 154: Copying the SSH User's Public Key	266
Figure 155: Showing the SSH User's Public Key	267
Figure 156: Showing TCAM Utilization	269
Figure 157: Creating an ACL	271
Figure 158: Showing a List of ACLs	271
Figure 159: Configuring a Standard IPv4 ACL	273
Figure 160: Configuring an Extended IPv4 ACL	275
Figure 161: Configuring a Standard IPv6 ACL	276
Figure 162: Configuring an Extended IPv6 ACL	278
Figure 163: Configuring a MAC ACL	280
Figure 164: Configuring a ARP ACL	282
Figure 165: Binding a Port to an ACL	283
Figure 166: Configuring ACL Mirroring	284
Figure 167: Showing the VLANs to Mirror	284
Figure 168: Showing ACL Statistics	286
Figure 169: Configuring Global Settings for ARP Inspection	289

Figure 170: Configuring VLAN Settings for ARP Inspection	290
Figure 171: Configuring Interface Settings for ARP Inspection	291
Figure 172: Displaying Statistics for ARP Inspection	293
Figure 173: Displaying the ARP Inspection Log	294
Figure 174: Creating an IP Address Filter for Management Access	295
Figure 175: Showing IP Addresses Authorized for Management Access	295
Figure 176: Configuring Port Security	298
Figure 177: Configuring Port Security	299
Figure 178: Configuring Global Settings for 802.1X Port Authentication	300
Figure 179: Configuring Interface Settings for 802.1X Port Authenticator	304
Figure 180: Showing Statistics for 802.1X Port Authenticator	306
Figure 181: Setting the Filter Type for IPv4 Source Guard	309
Figure 182: Configuring Static Bindings for IPv4 Source Guard	311
Figure 183: Displaying Static Bindings for IP Source Guard	311
Figure 184: Showing the IPv4 Source Guard Binding Table	313
Figure 185: Setting the Filter Type for IPv6 Source Guard	315
Figure 186: Configuring Static Bindings for IPv6 Source Guard	317
Figure 187: Displaying Static Bindings for IPv6 Source Guard	318
Figure 188: Showing the IPv6 Source Guard Binding Table	319
Figure 189: Configuring Global Settings for DHCP Snooping	323
Figure 190: Configuring DHCP Snooping on a VLAN	324
Figure 191: Configuring the Port Mode for DHCP Snooping	325
Figure 192: Displaying the Binding Table for DHCP Snooping	326
Figure 193: Configuring Settings for System Memory Logs	329
Figure 194: Showing Error Messages Logged to System Memory	330
Figure 195: Configuring Settings for Remote Logging of Error Messages	331
Figure 196: Configuring LLDP Timing Attributes	333
Figure 197: Configuring LLDP Interface Attributes	337
Figure 198: Configuring the Civic Address for an LLDP Interface	339
Figure 199: Showing the Civic Address for an LLDP Interface	339
Figure 200: Displaying Local Device Information for LLDP (General)	342
Figure 201: Displaying Local Device Information for LLDP (Port)	342
Figure 202: Displaying Local Device Information for LLDP (Port Details)	343
Figure 203: Displaying Remote Device Information for LLDP (Port)	349
Figure 204: Displaying Remote Device Information for LLDP (Port Details)	350

Figure 205: Displaying Remote Device Information for LLDP (End Node)	351
Figure 206: Displaying LLDP Device Statistics (General)	352
Figure 207: Displaying LLDP Device Statistics (Port)	353
Figure 208: Configuring Global Settings for SNMP	356
Figure 209: Configuring the Local Engine ID for SNMP	357
Figure 210: Configuring a Remote Engine ID for SNMP	358
Figure 211: Showing Remote Engine IDs for SNMP	358
Figure 212: Creating an SNMP View	359
Figure 213: Showing SNMP Views	360
Figure 214: Adding an OID Subtree to an SNMP View	360
Figure 215: Showing the OID Subtree Configured for SNMP Views	361
Figure 216: Creating an SNMP Group	365
Figure 217: Showing SNMP Groups	365
Figure 218: Setting Community Access Strings	366
Figure 219: Showing Community Access Strings	367
Figure 220: Configuring Local SNMPv3 Users	368
Figure 221: Showing Local SNMPv3 Users	369
Figure 222: Changing a Local SNMPv3 User Group	369
Figure 223: Configuring Remote SNMPv3 Users	371
Figure 224: Showing Remote SNMPv3 Users	372
Figure 225: Configuring Trap Managers (SNMPv1)	375
Figure 226: Configuring Trap Managers (SNMPv2c)	375
Figure 227: Configuring Trap Managers (SNMPv3)	376
Figure 228: Showing Trap Managers	376
Figure 229: Creating SNMP Notification Logs	378
Figure 230: Showing SNMP Notification Logs	378
Figure 231: Showing SNMP Statistics	380
Figure 232: Configuring an RMON Alarm	382
Figure 233: Showing Configured RMON Alarms	383
Figure 234: Configuring an RMON Event	384
Figure 235: Showing Configured RMON Events	385
Figure 236: Configuring an RMON History Sample	387
Figure 237: Showing Configured RMON History Samples	387
Figure 238: Showing Collected RMON History Samples	388
Figure 239: Configuring an RMON Statistical Sample	389

Figure 240: Showing Configured RMON Statistical Samples	389
Figure 241: Showing Collected RMON Statistical Samples	390
Figure 242: Single CFM Maintenance Domain	392
Figure 243: Multiple CFM Maintenance Domains	392
Figure 244: Configuring Global Settings for CFM	397
Figure 245: Configuring Interfaces for CFM	398
Figure 246: Configuring Maintenance Domains	401
Figure 247: Showing Maintenance Domains	402
Figure 248: Configuring Detailed Settings for Maintenance Domains	402
Figure 249: Creating Maintenance Associations	406
Figure 250: Showing Maintenance Associations	406
Figure 251: Configuring Detailed Settings for Maintenance Associations	407
Figure 252: Configuring Maintenance End Points	408
Figure 253: Showing Maintenance End Points	409
Figure 254: Configuring Remote Maintenance End Points	410
Figure 255: Showing Remote Maintenance End Points	410
Figure 256: Transmitting Link Trace Messages	412
Figure 257: Transmitting Loopback Messages	414
Figure 258: Transmitting Delay-Measure Messages	416
Figure 259: Showing Information on Local MEPs	417
Figure 260: Showing Detailed Information on Local MEPs	419
Figure 261: Showing Information on Local MIPs	420
Figure 262: Showing Information on Remote MEPs	421
Figure 263: Showing Detailed Information on Remote MEPs	423
Figure 264: Showing the Link Trace Cache	424
Figure 265: Showing Settings for the Fault Notification Generator	425
Figure 266: Showing Continuity Check Errors	426
Figure 267: Configuring UDLD Protocol Intervals	428
Figure 268: Configuring UDLD Interface Settings	430
Figure 269: Displaying UDLD Neighbor Information	431
Figure 270: Multicast Filtering Concept	433
Figure 271: IGMP Protocol	435
Figure 272: Configuring General Settings for IGMP Snooping	440
Figure 273: Configuring a Static Interface for an IPv4 Multicast Router	442
Figure 274: Showing Static Interfaces Attached an IPv4 Multicast Router	442

Figure 275: Showing Current Interfaces Attached an IPv4 Multicast Router	443
Figure 276: Assigning an Interface to an IPv4 Multicast Service	444
Figure 277: Showing Static Interfaces Assigned to an IPv4 Multicast Service	444
Figure 278: Showing Current Interfaces Attached a Multicast Router	445
Figure 279: Configuring IGMP Snooping on a VLAN	450
Figure 280: Showing Interface Settings for IGMP Snooping	451
Figure 281: Dropping IGMP Query Packets	452
Figure 282: Showing Multicast Groups Learned by IGMP Snooping	453
Figure 283: Displaying IGMP Snooping Statistics – Query	455
Figure 284: Displaying IGMP Snooping Statistics – VLAN	456
Figure 285: Displaying IGMP Snooping Statistics – Port	456
Figure 286: Enabling IGMP Filtering and Throttling	458
Figure 287: Creating an IGMP Filtering Profile	459
Figure 288: Showing the IGMP Filtering Profiles Created	459
Figure 289: Adding Multicast Groups to an IGMP Filtering Profile	460
Figure 290: Showing the Groups Assigned to an IGMP Filtering Profile	460
Figure 291: Configuring IGMP Filtering and Throttling Interface Settings	462
Figure 292: Configuring General Settings for MLD Snooping	464
Figure 293: Configuring Immediate Leave for MLD Snooping	465
Figure 294: Configuring a Static Interface for an IPv6 Multicast Router	466
Figure 295: Showing Static Interfaces Attached an IPv6 Multicast Router	466
Figure 296: Showing Current Interfaces Attached an IPv6 Multicast Router	466
Figure 297: Assigning an Interface to an IPv6 Multicast Service	468
Figure 298: Showing Static Interfaces Assigned to an IPv6 Multicast Service	468
Figure 299: Showing Current Interfaces Assigned to an IPv6 Multicast Service	469
Figure 300: Showing IPv6 Multicast Services and Corresponding Sources	470
Figure 301: IGMP Proxy Routing	471
Figure 302: Configuring IGMP Proxy Routing	473
Figure 303: Configuring IGMP Interface Settings	476
Figure 304: Configuring Static IGMP Groups	477
Figure 305: Showing Static IGMP Groups	477
Figure 306: Displaying Multicast Groups Learned from IGMP (Information)	480
Figure 307: Displaying Multicast Groups Learned from IGMP (update later)	480
Figure 308: Configuring a Static IPv4 Address	483
Figure 309: Configuring a Dynamic IPv4 Address	484

Figure 310: Showing the IPv4 Address Configured for an Interface	485
Figure 311: Configuring the IPv6 Default Gateway	486
Figure 312: Configuring General Settings for an IPv6 Interface	490
Figure 313: Configuring RA Guard for an IPv6 Interface	491
Figure 314: Configuring an IPv6 Address	493
Figure 315: Showing Configured IPv6 Addresses	495
Figure 316: Showing IPv6 Neighbors	496
Figure 317: Showing IPv6 Statistics (IPv6)	501
Figure 318: Showing IPv6 Statistics (ICMPv6)	502
Figure 319: Showing IPv6 Statistics (UDP)	502
Figure 320: Showing Reported MTU Values	503
Figure 321: Configuring General Settings for DNS	506
Figure 322: Configuring a List of Domain Names for DNS	507
Figure 323: Showing the List of Domain Names for DNS	507
Figure 324: Configuring a List of Name Servers for DNS	508
Figure 325: Showing the List of Name Servers for DNS	509
Figure 326: Configuring Static Entries in the DNS Table	509
Figure 327: Showing Static Entries in the DNS Table	510
Figure 328: Showing Entries in the DNS Cache	511
Figure 329: Specifying a DHCP Client Identifier	513
Figure 330: Layer 3 DHCP Relay Service	513
Figure 331: Configuring DHCP Relay Service	514
Figure 332: Virtual Interfaces and Layer 3 Routing	516
Figure 333: Pinging a Network Device	520
Figure 334: Tracing the Route to a Network Device	521
Figure 335: Configuring ARP Timeout	523
Figure 336: Configuring Static ARP Entries	524
Figure 337: Displaying Static ARP Entries	525
Figure 338: Displaying ARP Entries	525
Figure 339: Displaying ARP Statistics	526
Figure 340: Configuring Static Routes	527
Figure 341: Displaying Static Routes	528
Figure 342: Displaying the Routing Table	529
Figure 343: Setting the Maximum ECMP Number	531
Figure 344: Master Virtual Router with Backup Routers	533

Figure 345: Several Virtual Master Routers Using Backup Routers	534
Figure 346: Several Virtual Master Routers Configured for Mutual Backup and Load Sharing	534
Figure 347: Configuring the VRRP Group ID	538
Figure 348: Showing Configured VRRP Groups	538
Figure 349: Setting the Virtual Router Address for a VRRP Group	539
Figure 350: Showing the Virtual Addresses Assigned to VRRP Groups	539
Figure 351: Configuring Detailed Settings for a VRRP Group	540
Figure 352: Showing Counters for Errors Found in VRRP Packets	541
Figure 353: Showing Counters for Errors Found in a VRRP Group	542
Figure 354: Configuring RIP	544
Figure 355: Configuring General Settings for RIP	548
Figure 356: Clearing Entries from the Routing Table	549
Figure 357: Adding Network Interfaces to RIP	550
Figure 358: Showing Network Interfaces Using RIP	551
Figure 359: Specifying a Passive RIP Interface	551
Figure 360: Showing Passive RIP Interfaces	552
Figure 361: Specifying a Static RIP Neighbor	553
Figure 362: Showing Static RIP Neighbors	553
Figure 363: Redistributing External Routes into RIP	554
Figure 364: Showing External Routes Redistributed into RIP	555
Figure 365: Setting the Distance Assigned to External Routes	556
Figure 366: Showing the Distance Assigned to External Routes	556
Figure 367: Configuring a Network Interface for RIP	559
Figure 368: Showing RIP Network Interface Settings	560
Figure 369: Showing RIP Interface Settings	560
Figure 370: Showing RIP Peer Information	561
Figure 371: Resetting RIP Statistics	562
Figure 372: Configuring OSPF	562
Figure 373: OSPF Areas	564
Figure 374: Defining OSPF Network Areas Based on Addresses	565
Figure 375: Showing OSPF Network Areas	566
Figure 376: Showing OSPF Process Identifiers	566
Figure 377: AS Boundary Router	568
Figure 378: Configure General Settings for OSPF	569
Figure 379: Showing General Settings for OSPF	570

Figure 380: Adding an NSSA or Stub	571
Figure 381: Showing NSSAs or Stubs	572
Figure 382: OSPF NSSA	572
Figure 383: Configuring Protocol Settings for an NSSA	575
Figure 384: OSPF Stub Area	575
Figure 385: Configuring Protocol Settings for a Stub	577
Figure 386: Displaying Information on NSSA and Stub Areas	578
Figure 387: Route Summarization for ABRs	578
Figure 388: Configuring Route Summaries for an Area Range	579
Figure 389: Showing Configured Route Summaries	580
Figure 390: Redistributing External Routes	580
Figure 391: Importing External Routes	582
Figure 392: Showing Imported External Route Types	582
Figure 393: Summarizing External Routes	583
Figure 394: Showing Summary Addresses for External Routes	584
Figure 395: Configuring Settings for All Interfaces Assigned to a VLAN	587
Figure 396: Configuring Settings for a Specific Area Assigned to a VLAN	588
Figure 397: Showing OSPF Interfaces	589
Figure 398: Showing MD5 Authentication Keys	589
Figure 399: OSPF Virtual Link	590
Figure 400: Adding a Virtual Link	591
Figure 401: Showing Virtual Links	591
Figure 402: Configuring Detailed Settings for a Virtual Link	592
Figure 403: Showing MD5 Authentication Keys	592
Figure 404: Displaying Information in the Link State Database	594
Figure 405: Displaying Neighbor Routers Stored in the Link State Database	596
Figure 406: Specifying a Passive OSPF Interface	597
Figure 407: Showing Passive OSPF Interfaces	597
Figure 408: Enabling IPv4 Multicast Routing	602
Figure 409: Enabling IPv6 Multicast Routing	603
Figure 410: Displaying the IPv4 Multicast Routing Table	605
Figure 411: Displaying Detailed Entries from IPv4 Multicast Routing Table	606
Figure 412: Displaying the IPv6 Multicast Routing Table	606
Figure 413: Displaying Detailed Entries from IPv6 Multicast Routing Table	607
Figure 414: Enabling PIM Multicast Routing	608

Figure 415: Configuring PIM Interface Settings (Dense Mode)	612
Figure 416: Configuring PIM Interface Settings (Sparse Mode)	613
Figure 417: Showing PIM Neighbors	613
Figure 418: Configuring Global Settings for PIM-SM	615
Figure 419: Configuring a PIM-SM BSR Candidate	616
Figure 420: Configuring a PIM Static Rendezvous Point	618
Figure 421: Showing PIM Static Rendezvous Points	618
Figure 422: Configuring a PIM RP Candidate	620
Figure 423: Showing Settings for a PIM RP Candidate	620
Figure 424: Showing Information About the PIM BSR	622
Figure 425: Showing PIM RP Mapping	623
Figure 426: Enabling PIMv6 Multicast Routing	623
Figure 427: Configuring PIMv6 Interface Settings (Dense Mode)	628
Figure 428: Configuring PIMv6 Interface Settings (Sparse Mode)	629
Figure 429: Showing PIMv6 Neighbors	630
Figure 430: Configuring Global Settings for PIM6-SM	631
Figure 431: Configuring a PIM6-SM BSR Candidate	633
Figure 432: Configuring a PIM6 Static Rendezvous Point	634
Figure 433: Showing PIM6 Static Rendezvous Points	634
Figure 434: Configuring a PIM6 RP Candidate	636
Figure 435: Showing Settings for a PIM6 RP Candidate	637
Figure 436: Showing Information About the PIM6 BSR	638
Figure 437: Showing PIM6 RP Mapping	639

Tables

Table 1: Revision History	4
Table 2: Key Features	37
Table 3: System Defaults	44
Table 4: Web Page Configuration Buttons	51
Table 5: Switch Main Menu	52
Table 6: Port Statistics	114
Table 7: LACP Port Counters	135
Table 8: LACP Internal Configuration Information	136
Table 9: LACP Remote Device Configuration Information	138
Table 10: Traffic Segmentation Forwarding	142
Table 11: Recommended STA Path Cost Range	182
Table 12: Default STA Path Costs	183
Table 13: IEEE 802.1p Egress Queue Priority Mapping	203
Table 14: CoS Priority Levels	204
Table 15: Mapping Internal Per-hop Behavior to Hardware Queues	204
Table 16: Default Mapping of DSCP Values to Internal PHB/Drop Values	208
Table 17: Default Mapping of CoS/CFI to Internal PHB/Drop Precedence	210
Table 18: Mapping Internal PHB/Drop Precedence to CoS/CFI Values	213
Table 19: Mapping IP Precedence	214
Table 20: Default Mapping of IP Precedence to Internal PHB/Drop Values	215
Table 21: Dynamic QoS Profiles	247
Table 22: HTTPS System Support	256
Table 23: ARP Inspection Statistics	292
Table 24: ARP Inspection Log	293
Table 25: 802.1X Statistics	305
Table 26: Logging Levels	328
Table 27: LLDP MED Location CA Types	337
Table 28: Chassis ID Subtype	340
Table 29: System Capabilities	340

Table 30: Port ID Subtype	341
Table 31: Remote Port Auto-Negotiation Advertised Capability	345
Table 32: SNMPv3 Security Models and Levels	354
Table 33: Supported Notification Messages	362
Table 34: Remote MEP Priority Levels	400
Table 35: MEP Defect Descriptions	400
Table 36: Show IPv6 Neighbors - display description	495
Table 37: Show IPv6 Statistics - display description	497
Table 38: Show MTU - display description	503
Table 39: Options 60, 66 and 67 Statements	511
Table 40: Options 55 and 124 Statements	512
Table 41: Address Resolution Protocol	522
Table 42: ARP Statistics	525
Table 43: VRRP Group Statistics	541
Table 44: OSPF System Information	569
Table 45: Troubleshooting Chart	649

Section I

Getting Started

This section provides an overview of the switch, and introduces some basic concepts about network switches. It also describes the basic settings required to access the management interface.

This section includes these chapters:

- ◆ ["Introduction" on page 37](#)

1

Introduction

This switch provides a broad range of features for Layer 2 switching and Layer 3 routing. It includes a management agent that allows you to configure the features listed in this manual. The default configuration can be used for most of the features provided by this switch. However, there are many options that you should configure to maximize the switch's performance for your particular network environment.

Key Features

Table 2: Key Features

Feature	Description
Configuration Backup and Restore	Using management station or FTP/TFTP server
Authentication	Console, Telnet, web – user name/password, RADIUS, TACACS+ SNMP v1/2c - Community strings SNMP version 3 – MD5 or SHA password Telnet – SSH Web – HTTPS
General Security Measures	IP Address Filtering Local and Remote User Accounts RADIUS Server Authentication Secure Shell
Access Control Lists	Supports up to 256 ACLs, up to 96 rules per ACL
DHCP	Client, Relay
DHCPv6	Client
DNS	Client service
Port Configuration	Speed, duplex mode and flow control
Port Trunking	Supports up to 27 trunks on the AOS5700-54X and up to 16 trunks on the AOS6700-32X – static or dynamic trunking (LACP)
Port Mirroring	28 sessions, one or more source ports to one analysis port
Congestion Control	Rate Limiting Throttling for broadcast, multicast, unknown unicast storms

Table 2: Key Features (Continued)

Feature	Description
Address Table	32K MAC addresses in forwarding table, 1K static MAC addresses; 8K entries in ARP cache, 256 static ARP entries; 512 static IP routes, 512 IP interfaces; 12K IPv4 entries in host table; 8K IPv4 entries in routing table; 6K IPv6 entries in host table; 4K IPv6 entries in routing table; 1K L2 IPv4 multicast groups; 1K L3 IPv4 multicast groups (shared with IPv6); 1K L3 IPv6 multicast groups (shared with IPv4)
IP Version 4 and 6	Supports IPv4 and IPv6 addressing, and management
IEEE 802.1D Bridge	Supports dynamic data switching and addresses learning
Store-and-Forward Switching	Supported to ensure wire-speed switching while eliminating bad frames
Spanning Tree Algorithm	Supports standard STP, Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Trees (MSTP)
Virtual LANs	Up to 4094 using IEEE 802.1Q, and port-based VLANs
Traffic Prioritization	Default port priority, traffic class map, queue scheduling, IP Precedence, or Differentiated Services Code Point (DSCP), and TCP/UDP Port
Quality of Service	Supports Differentiated Services (DiffServ)
Link Layer Discovery Protocol	Used to discover basic information about neighboring devices
Router Redundancy	Router backup is provided with the Virtual Router Redundancy Protocol (VRRP)
IP Routing	Open Shortest Path First (OSPFv2/v3*), Border Gateway Protocol (BGPv4)*, policy-based routing for BGP*, static routes, Equal-Cost Multipath Routing (ECMP)
ARP	Static and dynamic address configuration, proxy ARP
Multicast Filtering	Supports IGMP snooping and query for Layer 2
Multicast Routing	Static multicast routing

* These features are only available through the Command Line Interface

Description of Software Features

The switch provides a wide range of advanced performance enhancing features. Flow control eliminates the loss of packets due to bottlenecks caused by port saturation. Broadcast storm suppression prevents broadcast traffic storms from engulfing the network. Untagged (port-based), tagged, and protocol-based VLANs, plus support for automatic GVRP VLAN registration provide traffic security and efficient use of network bandwidth. CoS priority queueing ensures the minimum delay for moving real-time multimedia data across the network. While multicast filtering and routing provides support for real-time network applications.

Some of the management features are briefly described below.

Configuration Backup and Restore You can save the current configuration settings to a file on the management station (using the web interface) or an FTP/TFTP server (using the web or console interface), and later download this file to restore the switch configuration settings.

Authentication This switch authenticates management access via the console port, Telnet, or a web browser. User names and passwords can be configured locally or can be verified via a remote authentication server (i.e., RADIUS or TACACS+).

Other authentication options include HTTPS for secure management access via the web, SSH for secure management access over a Telnet-equivalent connection, SNMP Version 3, IP address filtering for SNMP/Telnet/web management access. MAC address filtering and IP source guard also provide authenticated port access. Access Control Lists

ACLs provide packet filtering for IP frames (based on address, protocol, TCP/UDP port number or TCP control code) or any frames (based on MAC address or Ethernet type). ACLs can be used to improve performance by blocking unnecessary network traffic or to implement security controls by restricting access to specific network resources or protocols.

DHCP DHCP Relay Option 82 controls the processing of Option 82 information in DHCP request packets relayed by this device.

Port Configuration You can manually configure the speed, duplex mode, and flow control used on specific ports, or use auto-negotiation to detect the connection settings used by the attached device. Use full-duplex mode on ports whenever possible to double the throughput of switch connections. Flow control should also be enabled to control network traffic during periods of congestion and prevent the loss of packets when port buffer thresholds are exceeded. The switch supports flow control based on the IEEE 802.3x standard (now incorporated in IEEE 802.3-2002).

Port Mirroring The switch can unobtrusively mirror traffic from any port to a monitor port. You can then attach a protocol analyzer or RMON probe to this port to perform traffic analysis and verify connection integrity.

Port Trunking Ports can be combined into an aggregate connection. Trunks can be manually set up or dynamically configured using Link Aggregation Control Protocol (LACP – IEEE 802.3-2005). The additional ports dramatically increase the throughput across any connection, and provide redundancy by taking over the load if a port in the trunk should fail. The switch supports up to 27/16 trunks on the AOS5700-54X and AOS6700-32X respectively.

Storm Control Broadcast, multicast and unknown unicast storm suppression prevents traffic from overwhelming the network. When enabled on a port, the level of traffic passing through the port is restricted. If traffic rises above a pre-defined threshold, it will be throttled until the level falls back beneath the threshold.

Static MAC Addresses A static address can be assigned to a specific interface on this switch. Static addresses are bound to the assigned interface and will not be moved. When a static address is seen on another interface, the address will be ignored and will not be written to the address table. Static addresses can be used to provide network security by restricting access for a known host to a specific port.

IP Address Filtering Access to insecure ports can be controlled using DHCP Snooping which filters ingress traffic based on static IP addresses and addresses stored in the DHCP Snooping table. Traffic can also be restricted to specific source IP addresses or source IP/MAC address pairs based on static entries or entries stored in the DHCP Snooping table.

IEEE 802.1D Bridge The switch supports IEEE 802.1D transparent bridging. The address table facilitates data switching by learning addresses, and then filtering or forwarding traffic based on this information. The address table supports up to 32K addresses.

Store-and-Forward Switching The switch copies each frame into its memory before forwarding them to another port. This ensures that all frames are a standard Ethernet size and have been verified for accuracy with the cyclic redundancy check (CRC). This prevents bad frames from entering the network and wasting bandwidth.

To avoid dropping frames on congested ports, the switch provides 3 Mbits for frame buffering. This buffer can queue packets awaiting transmission on congested networks.

Spanning Tree Algorithm The switch supports these spanning tree protocols:

- ◆ Spanning Tree Protocol (STP, IEEE 802.1D) – This protocol provides loop detection. When there are multiple physical paths between segments, this protocol will choose a single path and disable all others to ensure that only one route exists between any two stations on the network. This prevents the creation of network loops. However, if the chosen path should fail for any reason, an alternate path will be activated to maintain the connection.
- ◆ Rapid Spanning Tree Protocol (RSTP, IEEE 802.1w) – This protocol reduces the convergence time for network topology changes to about 3 to 5 seconds, compared to 30 seconds or more for the older IEEE 802.1D STP standard. It is intended as a complete replacement for STP, but can still interoperate with switches running the older standard by automatically reconfiguring ports to

STP-compliant mode if they detect STP protocol messages from attached devices.

- ◆ Multiple Spanning Tree Protocol (MSTP, IEEE 802.1s) – This protocol is a direct extension of RSTP. It can provide an independent spanning tree for different VLANs. It simplifies network management, provides for even faster convergence than RSTP by limiting the size of each region, and prevents VLAN members from being segmented from the rest of the group (as sometimes occurs with IEEE 802.1D STP).

Virtual LANs The switch supports up to 4094 VLANs. A Virtual LAN is a collection of network nodes that share the same collision domain regardless of their physical location or connection point in the network. The switch supports tagged VLANs based on the IEEE 802.1Q standard. Members of VLAN groups can be manually assigned to a specific set of VLANs. This allows the switch to restrict traffic to the VLAN groups to which a user has been assigned. By segmenting your network into VLANs, you can:

- ◆ Eliminate broadcast storms which severely degrade performance in a flat network.
- ◆ Simplify network management for node changes/moves by remotely configuring VLAN membership for any port, rather than having to manually change the network connection.
- ◆ Provide data security by restricting all traffic to the originating VLAN, except where a connection is explicitly defined via the switch's routing service.

Traffic Prioritization This switch prioritizes each packet based on the required level of service, using eight priority queues with strict priority, Weighted Round Robin (WRR), or a combination of strict and weighted queuing. It uses IEEE 802.1p and 802.1Q tags to prioritize incoming traffic based on input from the end-station application. These functions can be used to provide independent priorities for delay-sensitive data and best-effort data.

This switch also supports several common methods of prioritizing layer 3/4 traffic to meet application requirements. Traffic can be prioritized based on the priority bits in the IP frame's Type of Service (ToS) octet using DSCP, or IP Precedence or TCP/UDP port numbers. When these services are enabled, the priorities are mapped to a Class of Service value by the switch, and the traffic then sent to the corresponding output queue.

Quality of Service Differentiated Services (DiffServ) provides policy-based management mechanisms used for prioritizing network resources to meet the requirements of specific traffic types on a per-hop basis. Each packet is classified upon entry into the network based on access lists, IP Precedence or DSCP values, or VLAN lists. Using access lists allows you select traffic based on Layer 2, Layer 3, or Layer 4 information contained

in each packet. Based on network policies, different kinds of traffic can be marked for different kinds of forwarding.

IP Routing The switch provides Layer 3 IP routing. To maintain a high rate of throughput, the switch forwards all traffic passing within the same segment, and routes only traffic that passes between different subnetworks. The wire-speed routing provided by this switch lets you easily link network segments or VLANs together without having to deal with the bottlenecks or configuration hassles normally associated with conventional routers.

Routing for unicast traffic is supported with static routing, Open Shortest Path First (OSPF) protocol, and Border Gateway Protocol (BGP).

Static Routing – Traffic is automatically routed between any IP interfaces configured on the switch. Routing to statically configured hosts or subnet addresses is provided based on next-hop entries specified in the static routing table.

OSPF – This approach uses a link state routing protocol to generate a shortest-path tree, then builds up its routing table based on this tree. OSPF produces a more stable network because the participating routers act on network changes predictably and simultaneously, converging on the best route more quickly than RIP.

BGP – This protocol uses a path vector approach to connect autonomous systems (AS) on the Internet. BGP maintains a table of IP network prefixes which designate network reachability among autonomous systems based the path of ASs to the destination, and next hop information. It makes routing decisions based on path, network policies and/or rule sets. For this reason, it is more appropriately termed a reachability protocol rather than a routing protocol.

Policy-based Routing for BGP – The next-hop behavior for ingress IP traffic can be determined based on matching criteria.

Equal-cost Multipath Load Balancing When multiple paths to the same destination and with the same path cost are found in the routing table, the Equal-cost Multipath (ECMP) algorithm first checks if the cost is lower than that of any other routing entries. If the cost is the lowest in the table, the switch will use up to eight paths having the lowest path cost to balance traffic forwarded to the destination. ECMP uses either equal-cost unicast multipaths manually configured in the static routing table, or equal-cost multipaths dynamically detected by the Open Shortest Path Algorithm (OSPF). In other words, it uses either static or unicast routing entries, not both.

Router Redundancy Virtual Router Redundancy Protocol (VRRP) uses a virtual IP address to support a primary router and multiple backup routers. The backups can be configured to take over the workload if the master fails or to load share the traffic. The primary goal of

this protocol is to allow a host device which has been configured with a fixed gateway to maintain network connectivity in case the primary gateway goes down.

Address Resolution Protocol

The switch uses ARP and Proxy ARP to convert between IP addresses and MAC (hardware) addresses. This switch supports conventional ARP, which locates the MAC address corresponding to a given IP address. This allows the switch to use IP addresses for routing decisions and the corresponding MAC addresses to forward packets from one hop to the next. Either static or dynamic entries can be configured in the ARP cache.

Proxy ARP allows hosts that do not support routing to determine the MAC address of a device on another network or subnet. When a host sends an ARP request for a remote network, the switch checks to see if it has the best route. If it does, it sends its own MAC address to the host. The host then sends traffic for the remote destination via the switch, which uses its own routing table to reach the destination on the other network.

Multicast Filtering

Specific multicast traffic can be assigned to its own VLAN to ensure that it does not interfere with normal network traffic and to guarantee real-time delivery by setting the required priority level for the designated VLAN. The switch uses IGMP Snooping and Query for IPv4.

Link Layer Discovery Protocol

LLDP is used to discover basic information about neighboring devices within the local broadcast domain. LLDP is a Layer 2 protocol that advertises information about the sending device and collects information gathered from neighboring network nodes it discovers.

Advertised information is represented in Type Length Value (TLV) format according to the IEEE 802.1ab standard, and can include details such as device identification, capabilities and configuration settings. Media Endpoint Discovery (LLDP-MED) is an extension of LLDP intended for managing endpoint devices such as Voice over IP phones and network switches. The LLDP-MED TLVs advertise information such as network policy, power, inventory, and device location details. The LLDP and LLDP-MED information can be used by SNMP applications to simplify troubleshooting, enhance network management, and maintain an accurate network topology.

System Defaults

The switch's system defaults are provided in the configuration file "Factory_Default_Config.cfg." To reset the switch defaults, this file should be set as the startup configuration file.

The following table lists some of the basic system defaults.

Table 3: System Defaults

Function	Parameter	Default
Console Port Connection	Baud Rate	115200 bps
	Data bits	8
	Stop bits	1
	Parity	none
	Local Console Timeout	0 (disabled)
Authentication and Security Measures	Privileged Exec Level	Username "admin" Password "admin"
	Normal Exec Level	Username "guest" Password "guest"
	Enable Privileged Exec from Normal Exec Level	Password "super"
	RADIUS Authentication	Disabled
	TACACS+ Authentication	Disabled
	MAC Authentication	Disabled
	HTTPS	Enabled
	SSH	Disabled
	IP Filtering	Disabled
	Web Management	HTTP Server
HTTP Port Number		80
HTTP Secure Server		Enabled
HTTP Secure Server Port		443
SNMP	SNMP Agent	Enabled
	Community Strings	"public" (read only) "private" (read/write)
	Traps	Authentication traps: enabled Link-up-down events: enabled
	SNMP V3	View: defaultview Group: public (read only); private (read/write)

Table 3: System Defaults (Continued)

Function	Parameter	Default
Port Configuration	Admin Status	Enabled
	Auto-negotiation	Enabled
	Flow Control	Disabled
Port Trunking	Static Trunks	None
	LACP (all ports)	Disabled
Congestion Control	Storm Control	Broadcast: Enabled (500 packets/sec) Multicast: Disabled Unknown Unicast: Disabled
Address Table	Aging Time	300 seconds
Spanning Tree Algorithm	Status	Enabled, RSTP (Defaults: RSTP standard)
	Edge Ports	Disabled
LLDP	Status	Enabled
Virtual LANs	Default VLAN	1
	PVID	1
	Acceptable Frame Type	All
	Ingress Filtering	Disabled
	Switchport Mode (Egress Mode)	Hybrid
Traffic Prioritization	Ingress Port Priority	0
	Queue Mode	WRR
	Queue Weight	Queue: 0 1 2 3 4 5 6 7 Weight: 1 2 4 6 8 10 12 14
	Class of Service	Enabled
	IP Precedence Priority	Disabled
IP Settings	IP DSCP Priority	Disabled
	IP Port Priority	Disabled
	Management. VLAN	VLAN 1
	IP Address	DHCP assigned
	Subnet Mask	255.255.255.0
	Default Gateway	0.0.0.0
	DHCP	Client: Enabled
	DNS	Client/Proxy service: Disabled
	ARP	Enabled Cache Timeout: 20 minutes Proxy: Disabled

Table 3: System Defaults (Continued)

Function	Parameter	Default
Unicast Routing	OSPF	Disabled
	OSPFv3	Disabled
	BGPv4	Disabled
Multicast Routing	Static	Disabled
Router Redundancy	VRRP	Disabled
Multicast Filtering	IGMP Snooping (Layer 2)	Snooping: Enabled Querier: Disabled
System Log	Status	Enabled
	Messages Logged to RAM	Levels 0-7 (all)
	Messages Logged to Flash	Levels 0-3
SNTP	Clock Synchronization	Disabled

Section II

Web Configuration

This section describes the basic switch features, along with a detailed description of how to configure each feature via a web browser.

This section includes these chapters:

- ◆ "Using the Web Interface" on page 49
- ◆ "Basic Management Tasks" on page 71
- ◆ "Interface Configuration" on page 103
- ◆ "VLAN Configuration" on page 145
- ◆ "Address Table Settings" on page 163
- ◆ "Spanning Tree Algorithm" on page 173
- ◆ "Congestion Control" on page 195
- ◆ "Class of Service" on page 199
- ◆ "Quality of Service" on page 219
- ◆ "Security Measures" on page 235
- ◆ "Basic Administration Protocols" on page 327
- ◆ "Multicast Filtering" on page 433
- ◆ "IP Configuration" on page 481
- ◆ "IP Services" on page 505
- ◆ "General IP Routing" on page 515
- ◆ "Unicast Routing" on page 543

2

Using the Web Interface

This switch provides an embedded HTTP web agent. Using a web browser you can configure the switch and view statistics to monitor network activity. The web agent can be accessed by any computer on the network using a standard web browser (Internet Explorer 8, or Mozilla Firefox 37, Google Chrome 42, or later versions).



Note: You can also use the Command Line Interface (CLI) to manage the switch over a serial connection to the console port or via Telnet. For more information on using the CLI, refer to the *CLI Reference Guide*.

Connecting to the Web Interface

Prior to accessing the switch from a web browser, be sure you have first performed the following tasks:

1. Configure the switch with a valid IP address, subnet mask, and default gateway using an out-of-band serial connection or DHCP protocol. (See the *CLI Reference Guide*.)
2. Set user names and passwords using an out-of-band serial connection. Access to the web agent is controlled by the same user names and passwords as the onboard configuration program. (See the *CLI Reference Guide*.)
3. After you enter a user name and password, you will have access to the system configuration program.



Note: You are allowed three attempts to enter the correct password; on the third failed attempt the current connection is terminated.

Note: If you log into the web interface as guest (Normal Exec level), you can view the configuration settings or change the guest password. If you log in as “admin” (Privileged Exec level), you can change the settings on any page.

Note: If the path between your management station and this switch does not pass through any device that uses the Spanning Tree Algorithm, then you can set the switch port attached to your management station to fast forwarding (i.e., enable Admin Edge Port) to improve the switch’s response time to management commands issued through the web interface. See [“Configuring Interface Settings for STA” on page 181](#).

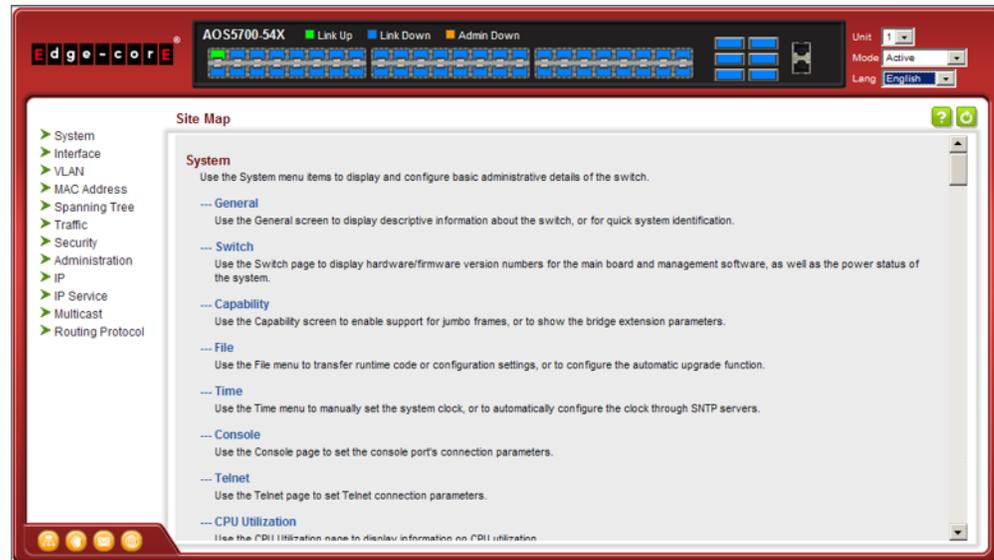
Note: Connection to the web interface is not supported for HTTPS using an IPv6 link local address.

Navigating the Web Browser Interface

To access the web-browser interface you must first enter a user name and password. The administrator has Read/Write access to all configuration parameters and statistics. The default user name and password for the administrator is “admin.”

Home Page When your web browser connects with the switch’s web agent, the home page is displayed as shown below. The home page displays the Main Menu on the left side of the screen and System Information on the right side. The Main Menu links are used to navigate to other menus, and display configuration parameters and statistics.

Figure 1: Home Page



Note: This manual covers the AS5700-54X 10G Ethernet switch and AS6700-32X 40G Ethernet switch, AS5700-54X and AS6700-32X are the bare metal switch names without any operating system installed. AOS5700-54X and AOS6700-32X are the same switches with the AOS operating system as described in this manual. Other than the difference in port types, there are no significant differences. Therefore nearly all of the screen display examples are based on the AOS5700-54X. The panel graphics for both switch types are shown on the following page.

Note: You can open a connection to the vendor’s web site by clicking on the Edge-Core logo.

Configuration Options Configurable parameters have a dialog box or a drop-down list. Once a configuration change has been made on a page, be sure to click on the Apply button to confirm the new setting. The following table summarizes the web page configuration buttons.

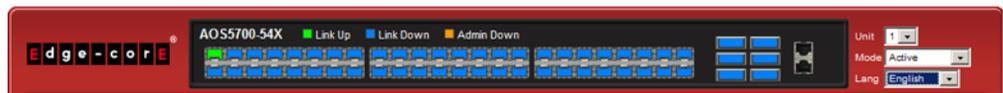
Table 4: Web Page Configuration Buttons

Button	Action
Apply	Sets specified values to the system.
Revert	Cancels specified values and restores current values prior to pressing "Apply."
	Displays help for the selected page.
	Refreshes the current page.
	Displays the site map.
	Logs out of the management interface.
	Sends mail to the vendor.
	Links to the vendor's web site.

Panel Display The web agent displays an image of the switch's ports. The Mode can be set to display different information for the ports, including Active (i.e., up or down), Duplex (i.e., half or full duplex), or Flow Control (i.e., with or without flow control).

Figure 2: Front Panel Indicators

AOS5700-54X



AOS6700-32X



Main Menu Using the onboard web agent, you can define system parameters, manage and control the switch, and all its ports, or monitor network conditions. The following table briefly describes the selections available from this program.

Table 5: Switch Main Menu

Menu	Description	Page
System		
General	Provides basic system description, including contact information	72
Switch	Shows the number of ports, hardware version, power status, and firmware version numbers	73
Capability	Enables support for jumbo frames; shows the bridge extension parameters	74, 75
File		77
Copy	Allows the transfer and copying files	77
Set Startup	Sets the startup file	80
Show	Shows the files stored in flash memory; allows deletion of files	80
Automatic Operation Code Upgrade	Automatically upgrades operation code if a newer version is found on the server	81
Time		85
Configure General		
Manual	Manually sets the current time	85
SNTP	Configures SNTP polling interval	86
NTP	Configures NTP authentication parameters	87
Configure Time Server	Configures a list of NTP or SNTP servers	88
Configure SNTP Server	Sets the IP address for SNTP time servers	88
Add NTP Server	Adds NTP time server and index of authentication key	89
Show NTP Server	Shows list of configured NTP time servers	89
Add NTP Authentication Key	Adds key index and corresponding MD5 key	90
Show NTP Authentication Key	Shows list of configured authentication keys	90
Configure Time Zone	Sets the local time zone for the system clock	92
Console	Sets console port connection parameters	93
Telnet	Sets Telnet connection parameters	95
CPU Utilization	Displays information on CPU utilization	97
Memory Status	Shows memory utilization parameters	98
Reset	Restarts the switch immediately, at a specified time, after a specified delay, or at a periodic interval	98

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Interface		103
Port		104
General		
Configure by Port List	Configures connection settings per port	104
Configure by Port Range	Configures connection settings for a range of ports	106
Show Information	Displays port connection status	107
Mirror		108
Add	Sets the source and target ports for mirroring	108
Show	Shows the configured mirror sessions	108
Statistics	Shows Interface, Etherlike, and RMON port statistics	114
Chart	Shows Interface, Etherlike, and RMON port statistics	114
History	Shows statistical history for the specified interfaces	119
Transceiver	Configures thresholds for alarm and warning messages for optical transceivers which support DDM	124
Trunk		
Static		127
Configure Trunk		127
Add	Creates a trunk, along with the first port member	127
Show	Shows the configured trunk identifiers	127
Add Member	Specifies ports to group into static trunks	127
Show Member	Shows the port members for the selected trunk	127
Configure General		127
Configure	Configures trunk connection settings	127
Show Information	Displays trunk connection settings	127
Dynamic		130
Configure Aggregator	Configures administration key and timeout for specific LACP groups	130
Configure Aggregation Port		127
Configure		127
General	Allows ports to dynamically join trunks	130
Actor	Configures parameters for link aggregation group members on the local side	130
Partner	Configures parameters for link aggregation group members on the remote side	130
Show Information		136
Counters	Displays statistics for LACP protocol messages	136

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Internal	Displays configuration settings and operational state for the local side of a link aggregation	137
Neighbors	Displays configuration settings and operational state for the remote side of a link aggregation	139
Configure Trunk		130
Configure	Configures connection settings	130
Show	Displays port connection status	130
Show Member	Shows the active members in a trunk	130
Statistics	Shows Interface, Etherlike, and RMON port statistics	114
Chart	Shows Interface, Etherlike, and RMON port statistics	114
Load Balance	Sets the load-distribution method among ports in aggregated links	140
History	Shows statistical history for the specified interfaces	119
Traffic Segmentation		142
Configure Global	Enables traffic segmentation globally	142
Configure Session	Configures the uplink and down-link ports for a segmented group of ports	143
Add	Assign the downlink and uplink ports to use in a segmented group	143
Show	Shows the assigned ports and direction (uplink/downlink)	143
VLAN	Virtual LAN	145
Static		
Add	Creates VLAN groups	147
Show	Displays configured VLAN groups	147
Modify	Configures group name and administrative status	147
Edit Member by VLAN	Specifies VLAN attributes per VLAN	150
Edit Member by Interface	Specifies VLAN attributes per interface	150
Edit Member by Interface Range	Specifies VLAN attributes per interface range	150
Tunnel	IEEE 802.1Q (QinQ) Tunneling	154
Configure Global	Sets tunnel mode for the switch	158
Configure Service	Sets a CVLAN to SPVLAN mapping entry	159
Configure Interface	Sets the tunnel mode for any participating interface	161
MAC Address		163
Learning Status	Enables MAC address learning on selected interfaces	163
Static		165
Add	Configures static entries in the address table	165
Show	Displays static entries in the address table	165

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Dynamic		
Configure Aging	Sets timeout for dynamically learned entries	167
Show Dynamic MAC	Displays dynamic entries in the address table	168
Clear Dynamic MAC	Removes any learned entries from the forwarding database and clears the transmit and receive counts for any static or system configured entries	169
MAC Notification		
Configure Global	Issues a trap when a dynamic MAC address is added or removed	170
Configure Interface	Enables MAC authentication traps on the current interface	170
Spanning Tree		173
STA		
Spanning Tree Algorithm		
Configure Global		
Configure	Configures global bridge settings for STP, RSTP and MSTP	175
Show Information	Displays STA values used for the bridge	180
Configure Interface		
Configure	Configures interface settings for STA	181
Show Inform at on	Displays interface settings for STA	186
MSTP		
Multiple Spanning Tree Algorithm		
Configure Global		
Add	Configures initial VLAN and priority for an MST instance	188
Show	Configures global settings for an MST instance	188
Modify	Configures the priority or an MST instance	188
Add Member	Adds VLAN members for an MST instance	188
Show Member	Adds or deletes VLAN members for an MST instance	188
Show Information	Displays MSTP values used for the bridge	
Configure Interface		192
Configure	Configures interface settings for an MST instance	192
Show Information	Displays interface settings for an MST instance	192
Traffic		
Rate Limit		
Rate Limit	Sets the input and output rate limits for a port	195
Storm Control		
Storm Control	Sets the broadcast storm threshold for each interface	196
Priority		
Default Priority		
Default Priority	Sets the default priority for each port or trunk	199
Queue		
Queue	Sets queue mode for the switch; sets the service weight for each queue that will use a weighted or hybrid mode	200
Trust Mode		
Trust Mode	Selects DSCP or CoS priority processing	206

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
DSCP to DSCP		207
Add	Maps DSCP values in incoming packets to per-hop behavior and drop precedence values for internal priority processing	207
Show	Shows the DSCP to DSCP mapping list	207
CoS to DSCP		210
Configure	Maps CoS/CFI values in incoming packets to per-hop behavior and drop precedence values for priority processing	210
Show	Shows the CoS to DSCP mapping list	210
DSCP to CoS		212
Add	Maps internal per-hop behavior and drop precedence value pairs to CoS values used in tagged egress packets on a Layer 2 interface	212
Show	Shows the DSCP to CoS mapping list	212
IP Precedence to DSCP		214
Add	Maps IP precedence values in incoming packets to per-hop behavior and drop precedence values for priority processing	214
Show	Shows the IP Precedence to DSCP mapping list	214
IP Port to DSCP		216
Add	Sets TCP/UDP port priority, defining the socket number and associated per-hop behavior and drop precedence	216
Show	Shows the IP Port to DSCP mapping list	216
PHB to Queue		203
Configure	Maps internal per-hop behavior values to hardware queues	203
Show	Shows the PHB to Queue mapping list	203
DiffServ		219
Configure Class		220
Add	Creates a class map for a type of traffic	220
Show	Shows configured class maps	220
Modify	Modifies the name of a class map	220
Add Rule	Configures the criteria used to classify ingress traffic	220
Show Rule	Shows the traffic classification rules for a class map	220
Configure Policy		224
Add	Creates a policy map to apply to multiple interfaces	224
Show	Shows configured policy maps	224
Modify	Modifies the name of a policy map	224
Add Rule	Sets the boundary parameters used for monitoring inbound traffic, and the action to take for conforming and non-conforming traffic	224
Show Rule	Shows the rules used to enforce bandwidth policing for a policy map	224

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Configure Interface	Applies a policy map to an ingress port	233
Security		235
AAA	Authentication, Authorization and Accounting	236
System Authentication	Configures authentication sequence – local, RADIUS, and TACACS	237
Server	Configures RADIUS and TACACS server message exchange settings	238
User Accounts		241
Add	Configures user names, passwords, and access levels	241
Show	Shows authorized users	241
Modify	Modifies user attributes	241
Web Authentication	Allows authentication and access to the network when 802.1X or Network Access authentication are infeasible or impractical	243
Configure Global	Configures general protocol settings	244
Configure Interface	Enables Web Authentication for individual ports	245
Network Access	MAC address-based network access authentication	246
Configure Global	Enables aging for authenticated MAC addresses, and sets the time period after which a connected MAC address must be reauthenticated	248
Configure Interface		249
General	Enables MAC authentication on a port; sets the maximum number of address that can be authenticated, the guest VLAN, dynamic VLAN and dynamic QoS	249
Link Detection	Configures detection of changes in link status, and the response (i.e., send trap or shut down port)	251
Configure MAC Filter		252
Add	Specifies MAC addresses exempt from authentication	252
Show	Shows the list of exempt MAC addresses	252
Show Information	Shows the authenticated MAC address list	254
HTTPS	Secure HTTP	255
Configure Global	Enables HTTPS, and specifies the UDP port to use	255
Copy Certificate	Replaces the default secure-site certificate	257
SSH	Secure Shell	258
Configure Global	Configures SSH server settings	261
Configure Host Key		262
Generate	Generates the host key pair (public and private)	262
Show	Displays RSA and DSA host keys; deletes host keys	262
Configure User Key		264
Copy	Imports user public keys from TFTP server	264
Show	Displays RSA and DSA user keys; deletes user keys	264

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
ACL	Access Control Lists	266
Configure ACL		269
Show TCAM	Shows utilization parameters for TCAM	268
Add	Adds an ACL based on IP or MAC address filtering	269
Show	Shows the name and type of configured ACLs	269
Add Rule	Configures packet filtering based on IP or MAC addresses and other packet attributes	269
Show Rule	Shows the rules specified for an ACL	269
Configure Interface		
Configure	Binds a port to the specified ACL	281
Add Mirror	Mirrors matching traffic to the specified port	282
Show Mirror	Shows ACLs mirrored to specified port	282
Show Hardware Counters	Shows statistics for ACL hardware counters	284
ARP Inspection		285
Configure General	Enables inspection globally, configures validation of additional address components, and sets the log rate for packet inspection	286
Configure VLAN	Enables ARP inspection on specified VLANs	288
Configure Interface	Sets the trust mode for ports, and sets the rate limit for packet inspection	289
Show Information		
Show Statistics	Displays statistics on the inspection process	291
Show Log	Shows the inspection log list	292
IP Filter		293
Add	Sets IP addresses of clients allowed management access via the web, SNMP, and Telnet	293
Show	Shows the addresses to be allowed management access	293
Port Security	Configures per port security, including status, response for security breach, and maximum allowed MAC addresses	295
Port Authentication	IEEE 802.1X	297
Configure Global	Enables authentication and EAPOL pass-through	299
Configure Interface	Sets authentication parameters for individual ports	300
Show Statistics	Displays protocol statistics for the selected port	304
IP Source Guard	Filters IPv4 traffic based on static entries in the IP Source Guard table, or dynamic entries in the DHCPv4 Snooping table	306
Port Configuration	Enables IP source guard, selects filter type per port, and sets maximum binding entries	306
Static Binding		308

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Configure ACL Table		308
Add	Adds a static addresses to the source-guard binding table	308
Show	Shows static addresses in the source-guard binding table	308
Configure MAC Table		308
Add	Adds a static addresses to the source-guard binding table	308
Show	Shows static addresses in the source-guard binding table	308
Dynamic Binding	Displays the source-guard binding table for a selected interface	311
IPv6 Source Guard	Filters IPv6 traffic based on static entries in the IP Source Guard table, or dynamic entries in the DHCP Snooping table	312
Port Configuration	Enables IPv6 source guard and selects filter type per port	312
Static Binding		315
Add	Adds a static addresses to the source-guard binding table	315
Show	Shows static addresses in the source-guard binding table	315
Dynamic Binding	Displays the source-guard binding table for a selected interface	317
Administration		327
Log		327
System		327
Configure Global	Stores error messages in local memory	327
Show System Logs	Shows logged error messages	327
Remote	Configures the logging of messages to a remote logging process	330
LLDP		331
Configure Global	Configures global LLDP timing parameters	332
Configure Interface	Sets the message transmission mode; enables SNMP notification; and sets the LLDP attributes to advertise	334
Show Local Device Information		339
General	Displays general information about the local device	339
Port/Trunk	Displays information about each interface	339
Show Remote Device Information		343
Port/Trunk	Displays information about a remote device connected to a port on this switch	343
Port/Trunk Details	Displays detailed information about a remote device connected to this switch	343
Show Device Statistics		351
General	Displays statistics for all connected remote devices	351
Port/Trunk	Displays statistics for remote devices on a selected port or trunk	351

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
SNMP	Simple Network Management Protocol	353
Configure Global	Enables SNMP agent status, and sets related trap functions	355
Configure Engine		356
Set Engine ID	Sets the SNMP v3 engine ID on this switch	356
Add Remote Engine	Sets the SNMP v3 engine ID for a remote device	357
Show Remote Engine	Shows configured engine ID for remote devices	357
Configure View		358
Add View	Adds an SNMP v3 view of the OID MIB	358
Show View	Shows configured SNMP v3 views	358
Add OID Subtree	Specifies a part of the subtree for the selected view	358
Show OID Subtree	Shows the subtrees assigned to each view	358
Configure Group		361
Add	Adds a group with access policies for assigned users	361
Show	Shows configured groups and access policies	361
Configure User		
Add Community	Configures community strings and access mode	366
Show Community	Shows community strings and access mode	366
Add SNMPv3 Local User	Configures SNMPv3 users on this switch	367
Show SNMPv3 Local User	Shows SNMPv3 users configured on this switch	367
Change SNMPv3 Local User Group	Assign a local user to a new group	367
Add SNMPv3 Remote User	Configures SNMPv3 users from a remote device	369
Show SNMPv3 Remote User	Shows SNMPv3 users set from a remote device	367
Configure Trap		372
Add	Configures trap managers to receive messages on key events that occur this switch	372
Show	Shows configured trap managers	372
Configure Notify Filter		376
Add	Creates an SNMP notification log	376
Show	Shows the configured notification logs	376
Show Statistics	Shows the status of SNMP communications	378
RMON	Remote Monitoring	380
Configure Global		
Add		
Alarm	Sets threshold bounds for a monitored variable	380
Event	Creates a response event for an alarm	383

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Show		
Alarm	Shows all configured alarms	380
Event	Shows all configured events	383
Configure Interface		
Add		
History	Periodically samples statistics on a physical interface	385
Statistics	Enables collection of statistics on a physical interface	388
Show		
History	Shows sampling parameters for each entry in the history group	385
Statistics	Shows sampling parameters for each entry in the statistics group	388
Show Details		
History	Shows sampled data for each entry in the history group	385
Statistics	Shows sampled data for each entry in the history group	388
CFM	Connectivity Fault Management	390
Configure Global	Configures global settings, including administrative status, cross-check start delay, link trace, and SNMP traps	394
Configure Interface	Configures administrative status on an interface	397
Configure MD	Configure Maintenance Domains	398
Add	Defines a portion of the network for which connectivity faults can be managed, identified by an MD index, maintenance level, and the MIP creation method	398
Configure Details	Configures the archive hold time and fault notification settings	398
Show	Shows list of configured maintenance domains	398
Configure MA	Configure Maintenance Associations	403
Add	Defines a unique CFM service instance, identified by its parent MD, the MA index, the VLAN assigned to the MA, and the MIP creation method	403
Configure Details	Configures detailed settings, including continuity check status and interval level, cross-check status, and alarm indication signal parameters	403
Show	Shows list of configured maintenance associations	403
Configure MEP	Configures Maintenance End Points	407
Add	Configures MEPs at the domain boundary to provide management access for each maintenance association	407
Show	Shows list of configured maintenance end points	407
Configure Remote MEP	Configures Remote Maintenance End Points	409
Add	Configures a static list of remote MEPs for comparison against the MEPs learned through continuity check messages	409
Show	Shows list of configured remote maintenance end points	409

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Transmit Link Trace	Sends link trace messages to isolate connectivity faults by tracing the path through a network to the designated target node	411
Transmit Loopback	Sends loopback messages to isolate connectivity faults by requesting a target node to echo the message back to the source	412
Transmit Delay Measure	Sends periodic delay-measure requests to a specified MEP within a maintenance association	414
Show Information		
Show Local MEP	Shows the MEPs configured on this device	523
Show Local MEP Details	Displays detailed CFM information about a specified local MEP in the continuity check database	524
Show Local MIP	Shows the MIPs on this device discovered by the CFM protocol	526
Show Remote MEP	Shows MEPs located on other devices which have been discovered through continuity check messages, or statically configured in the MEP database	527
Show Remote MEP Details	Displays detailed CFM information about a specified remote MEP in the continuity check database	528
Show Link Trace Cache	Shows information about link trace operations launched from this device	530
Show Fault Notification Generator	Displays configuration settings for the fault notification generator	532
Show Continuity Check Error	Displays CFM continuity check errors logged on this device	533
UDLD	UniDirectional Link Detection	427
Configure Global	Configures the message probe interval, detection interval, and recovery interval	427
Configure Interface	Enables UDLD and aggressive mode which reduces the shut-down delay after loss of bidirectional connectivity is detected	429
Show Information	Displays UDLD neighbor information, including neighbor state, expiration time, and protocol intervals	430
IP		
General		
Routing Interface		
Add Address	Configures an IP interface for a VLAN	481
Show Address	Shows the IP interfaces assigned to a VLAN	481
Ping	Sends ICMP echo request packets to another node on the network	519
Trace Route	Shows the route packets take to the specified destination	520
ARP	Address Resolution Protocol	522
Configure Static Address		523
Add	Statically maps a physical address to an IP address	523
Show	Shows the MAC to IP address static table	523

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Show Information		525
Dynamic Address	Shows dynamically learned entries in the IP routing table	525
Other Address	Shows internal addresses used by the switch	525
Statistics	Shows statistics on ARP requests sent and received	525
Routing		
Static Routes		526
Add	Configures static routing entries	526
Show	Shows static routing entries	526
Routing Table		528
Show Information	Shows all routing entries, including local, static and dynamic routes	528
Configure ECMP Number	Sets the maximum number of equal-cost paths to the same destination that can be installed in the routing table	529
VRRP	Virtual Router Redundancy Protocol	533
Configure Group ID		534
Add	Adds a VRRP group identifier to a VLAN	534
Show	Shows the VRRP group identifier list	534
Add IP Address	Sets a virtual interface address for a VRRP group	534
Show IP Address	Shows the virtual interface address assigned to a VRRP group	534
Configure Detail	Configure detailed settings, such as advertisement interval, preemption, priority, and authentication	534
Show Statistics		
Global Statistics	Displays global statistics for VRRP protocol packet errors	540
Group Statistics	Displays statistics for VRRP protocol events and errors on the specified VRRP group and interface	541
IPv6 Configuration		485
Configure Global	Sets an IPv6 default gateway for traffic with no known next hop	485
Configure Interface	Configures IPv6 interface address using auto-configuration or link-local address, and sets related protocol settings	486
Add IPv6 Address	Adds an global unicast, EUI-64, or link-local IPv6 address to an interface	491
Show IPv6 Address	Show the IPv6 addresses assigned to an interface	494
Show IPv6 Neighbor Cache	Displays information in the IPv6 neighbor discovery cache	495
Show Statistics		496
IPv6	Shows statistics about IPv6 traffic	496
ICMPv6	Shows statistics about ICMPv6 messages	496
UDP	Shows statistics about UDP messages	496

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Show MTU	Shows the maximum transmission unit (MTU) cache for destinations that have returned an ICMP packet-too-big message along with an acceptable MTU to this switch	503
IP Service		505
DNS	Domain Name Service	505
General		505
Configure Global	Enables DNS lookup; defines the default domain name appended to incomplete host names	505
Add Domain Name	Defines a list of domain names that can be appended to incomplete host names	506
Show Domain Names	Shows the configured domain name list	506
Add Name Server	Specifies IP address of name servers for dynamic lookup	508
Show Name Servers	Shows the name server address list	508
Static Host Table		509
Add	Configures static entries for domain name to address mapping	509
Show	Shows the list of static mapping entries	509
Modify	Modifies the static address mapped to the selected host name	509
Cache	Displays cache entries discovered by designated name servers	510
DHCP	Dynamic Host Configuration Protocol	
Client	Specifies the DHCP client identifier for an interface	511
Relay	Specifies DHCP relay servers	513
Snooping		318
Configure Global	Enables DHCP snooping globally, MAC-address verification, information option; and sets the information policy	320
Configure VLAN	Enables DHCP snooping on a VLAN	322
Configure Interface	Sets the trust mode for an interface	323
Show Information	Displays the DHCP Snooping binding information	324
Multicast		433
IGMP Snooping		435
General	Enables multicast filtering; configures parameters for IPv4 multicast snooping	437
Multicast Router		441
Add Static Multicast Router	Assigns ports that are attached to a neighboring multicast router	441
Show Static Multicast Router	Displays ports statically configured as attached to a neighboring multicast router	441
Show Current Multicast Router	Displays ports attached to a neighboring multicast router, either through static or dynamic configuration	441

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
IGMP Member		443
Add Static Member	Statically assigns multicast addresses to the selected VLAN	443
Show Static Member	Shows multicast addresses statically configured on the selected VLAN	443
Show Current Member	Shows multicast addresses associated with the selected VLAN, either through static or dynamic configuration	443
Interface		445
Configure VLAN	Configures IGMP snooping per VLAN interface	445
Configure Interface	Configures the interface to drop IGMP query packets or all multicast data packets	451
Show VLAN Information	Shows IGMP snooping settings per VLAN interface	445
Forwarding Entry	Displays the current multicast groups learned through IGMP Snooping	452
Filter		457
Configure General	Enables IGMP filtering for the switch	457
Configure Profile		458
Add	Adds IGMP filter profile; and sets access mode	458
Show	Shows configured IGMP filter profiles	458
Add Multicast Group Range	Assigns multicast groups to selected profile	458
Show Multicast Group Range	Shows multicast groups assigned to a profile	458
Configure Interface	Assigns IGMP filter profiles to port interfaces and sets throttling action	461
Statistics		453
Show Query Statistics	Shows statistics for query-related messages	453
Show VLAN Statistics	Shows statistics for protocol messages and number of active groups	453
Show Port Statistics	Shows statistics for protocol messages and number of active groups	453
Show Trunk Statistics	Shows statistics for protocol messages and number of active groups	453
MLD Snooping		462
General	Enables multicast filtering; configures parameters for IPv6 multicast snooping	462
Interface	Configures Immediate Leave status for a VLAN	464
Multicast Router		
Add Static Multicast Router	Assigns ports that are attached to a neighboring multicast router	465
Show Static Multicast Router	Displays ports statically configured as attached to a neighboring multicast router	465
Show Current Multicast Router	Displays ports attached to a neighboring multicast router, either through static or dynamic configuration	465
MLD Member		
Add Static Member	Statically assigns multicast addresses to the selected VLAN	467

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Show Static Member	Shows multicast addresses statically configured on the selected VLAN	467
Show Current Member	Shows multicast addresses associated with the selected VLAN, either through static or dynamic configuration	467
Group Information	Displays known multicast groups, member ports, the means by which each group was learned, and the corresponding source list	469
IGMP		470
Proxy	Configures IGMP proxy service for multicast routing	471
Interface	Configures Layer 3 IGMP settings for selected VLAN interface	474
Static Group		476
Add	Configures the router to be a static member of a multicast group on the specified VLAN interface	476
Show	Shows multicast group statically assigned to a VLAN interface	476
Group Information		478
Show Information	Shows the current multicast groups learned through IGMP for each VLAN	478
Show Details	Shows detailed information on each multicast group associated with a VLAN interface	478
Multicast Routing		599
General	Globally enables IPv4 multicast routing	602
Information		603
Show Summary	Shows each multicast route the switch has learned	603
Show Details	Shows additional information for each multicast route the switch has learned, including RP address, upstream router, and downstream interfaces	603
IPv6 Multicast Routing		599
General	Globally enables IPv6 multicast routing	602
Information		603
Show Summary	Shows each multicast route the switch has learned	603
Show Details	Shows additional information for each multicast route the switch has learned, including RP address, upstream router, and downstream interfaces	603
Routing Protocol		543
RIP		544
General		545
Configure	Enables or disables RIP, sets the global RIP attributes and timer values	545
Clear Route	Clears the specified route type or network interface from the routing table	548
Network		549
Add	Sets the network interfaces that will use RIP	549

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Show	Shows the network interfaces that will use RIP	549
Passive Interface		551
Add	Stops RIP broadcast and multicast messages from being sent on specified network interfaces	551
Show	Shows the configured passive interfaces	551
Neighbor Address		552
Add	Configures the router to directly exchange routing information with a static neighbor	552
Show	Shows adjacent hosts or interfaces configured as a neighboring router	552
Redistribute		553
Add	Imports external routing information from other routing domains (that is, protocols) into the autonomous system	553
Show	Shows the external routing information to be imported from other routing domains	553
Distance		555
Add	Defines an administrative distance for external routes learned from other routing protocols	555
Show	Shows the administrative distances assigned to external routes learned from other routing protocols	555
Interface		556
Add	Configures RIP parameters for each interface, including send and receive versions, authentication, and method of loopback prevention	556
Show	Shows the RIP parameters set for each interface	556
Modify	Modifies RIP parameters for an interface	556
Statistics		
Show Interface Information	Shows RIP settings, and statistics on RIP protocol messages	560
Show Peer Information	Displays information on neighboring RIP routers	561
Reset Statistics	Clears statistics for RIP protocol messages	561
OSPF	Open Shortest Path First (Version 2)	562
Network Area		563
Add	Defines OSPF area address, area ID, and process ID	563
Show	Shows configured areas	563
Show Process	Show configured processes	563
System		566
Configure	Configures the Router ID, global settings, and default information	566
Show	Shows LSA statistics, administrative status, ABR/ASBR, area count, and version number	569

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Area		
Configure Area		571
Add Area	Adds NSSA or stub	571
Show Area	Shows configured NSSA or stub	571
Configure NSSA Area	Configures settings for importing routes into or exporting routes out of not-so-stubby areas	572
Configure Stub Area	Configures default cost, and settings for importing routes into a stub	575
Show Information	Shows statistics for each area, including SPF startups, ABR/ASBR count, LSA count, and LSA checksum	577
Area Range		578
Add	Configures route summaries to advertise at an area boundary	578
Show	Shows route summaries advertised at an area boundary	578
Modify	Modifies route summaries advertised at an area boundary	578
Redistribute		580
Add	Redistributes routes from one routing domain to another	580
Show	Shows route types redistributed to another domain	580
Modify	Modifies configuration settings for redistributed routes	580
Summary Address		582
Add	Aggregates routes learned from other protocols for advertising into other autonomous systems	582
Show	Shows configured summary addresses	582
Interface		584
Show	Shows area ID and designated router settings for each interface	584
Configure by VLAN	Configures OSPF protocol settings and authentication for specified VLAN	584
Configure by Address	Configures OSPF protocol settings and authentication for specified interface address	584
Show MD5 Key	Shows MD5 key ID used for each area	584
Virtual Link		589
Add	Configures a virtual link through a transit area to the backbone	589
Show	Shows virtual links, neighbor address, and state	589
Configure Detailed Settings	Configures detailed protocol and authentication settings	589
Show MD5 Key	Shows the MD5 key ID used for each neighbor	589
Information		592
LSDB	Shows information about different OSPF Link State Advertisements (LSAs)	592
Neighbor	Shows information about each OSPF neighbor	595

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Passive Interface	Suppresses OSPF routing traffic on the specified interface	551
Add	Adds passive interface	551
Show	Shows passive interfaces	551
PIM		607
General	Enables PIM globally for the switch	607
Interface	Enables PIM per interface, and sets the mode to dense or sparse	608
Neighbor	Displays information neighboring PIM routers	613
SM		
Configure Global	Configures settings for register messages, and use of the SPT	614
BSR Candidate	Configures the switch as a BSR candidate	615
RP Address		617
Add	Sets a static address for an RP and the associated multicast group(s)	617
Show	Shows the static addresses configured for each RP and the associated multicast groups	617
RP Candidate		618
Add	Advertises the switch as an RP candidate to the BSR for the specified multicast groups	618
Show	Shows the multicast groups for which this switch is advertising itself as an RP candidate to the BSR	618
Show Information		
Show BSR Router	Displays information about the BSR	620
Show RP Mapping	Displays the active RPs and associated multicast routing entries	622
PIM6	PIM for IPv6	
General	Enables PIM globally for the switch	623
Interface	Enables PIM per interface, and sets the mode to dense or sparse	624
Neighbor	Displays information neighboring PIM routers	629
SM		
Configure Global	Configures settings for register messages, and use of the SPT	630
BSR Candidate	Configures the switch as a BSR candidate	631
RP Address		633
Add	Sets a static address for an RP and the associated multicast group(s)	633
Show	Shows the static addresses configured for each RP and the associated multicast groups	633
RP Candidate		635
Add	Advertises the switch as an RP candidate to the BSR for the specified multicast groups	635

Table 5: Switch Main Menu (Continued)

Menu	Description	Page
Show	Shows the multicast groups for which this switch is advertising itself as an RP candidate to the BSR	635
Show Information		
Show BSR Router	Displays information about the BSR	637
Show RP Mapping	Displays the active RPs and associated multicast routing entries	638

3

Basic Management Tasks

This chapter describes the following topics:

- ◆ [Displaying System Information](#) – Provides basic system description, including contact information.
- ◆ [Displaying Hardware/Software Versions](#) – Shows the hardware version, power status, and firmware versions
- ◆ [Configuring Support for Jumbo Frames](#) – Enables support for jumbo frames.
- ◆ [Displaying Bridge Extension Capabilities](#) – Shows the bridge extension parameters.
- ◆ [Managing System Files](#) – Describes how to upgrade operating software or configuration files, and set the system start-up files.
- ◆ [Setting the System Clock](#) – Sets the current time manually or through specified NTP or SNTP servers.
- ◆ [Configuring The Console Port](#) – Sets console port connection parameters.
- ◆ [Configuring Telnet Settings](#) – Sets Telnet connection parameters.
- ◆ [Displaying CPU Utilization](#) – Displays information on CPU utilization.
- ◆ [Displaying Memory Utilization](#) – Shows memory utilization parameters.

- ◆ **Resetting the System** – Restarts the switch immediately, at a specified time, after a specified delay, or at a periodic interval.

Displaying System Information

Use the System > General page to identify the system by displaying information such as the device name, location and contact information.

Parameters

These parameters are displayed:

- ◆ **System Description** – Brief description of device type.
- ◆ **System Object ID** – MIB II object ID for switch's network management subsystem.
- ◆ **System Up Time** – Length of time the management agent has been up.
- ◆ **System Name** – Name assigned to the switch system.
- ◆ **System Location** – Specifies the system location.
- ◆ **System Contact** – Administrator responsible for the system.

Web Interface

To configure general system information:

1. Click System, General.
2. Specify the system name, location, and contact information for the system administrator.
3. Click Apply.

Figure 3: System Information

The screenshot shows the 'System > General' configuration page. It displays the following information:

System Description	AOS5700-54X
System Object ID	1.3.6.1.4.1.259.12.1.2.101
System Up Time	0 days, 2 hours, 6 minutes, and 14. 24 seconds
System Name	<input type="text"/>
System Location	<input type="text"/>
System Contact	<input type="text"/>

At the bottom right, there are two buttons: 'Apply' and 'Revert'.

Displaying Hardware/Software Versions

Use the System > Switch page to display hardware/firmware version numbers for the main board and management software, as well as the power status of the system.

Parameters

The following parameters are displayed:

Main Board Information

- ◆ **Serial Number** – The serial number of the switch.
- ◆ **Number of Ports** – Number of built-in ports.
- ◆ **Hardware Version** – Hardware version of the main board.
- ◆ **Main Power Status** – Displays the status of the internal power supply.
- ◆ **Redundant Power Status** – Displays the status of the redundant power supply.

Management Software Information

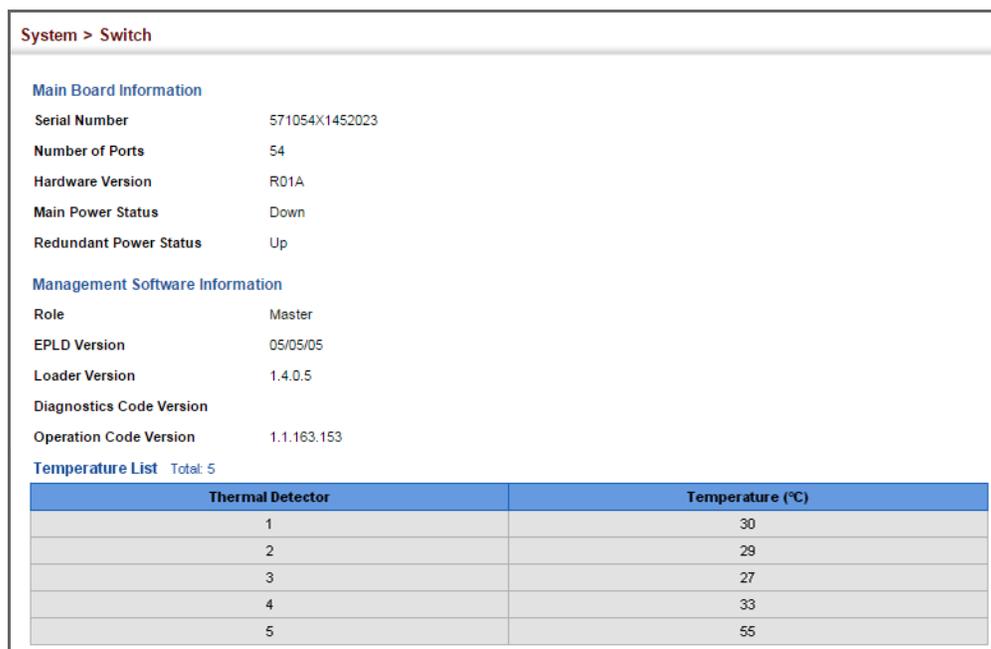
- ◆ **Role** – Shows that this switch is operating as Master or Slave.
- ◆ **EPLD Version** – Version number of EEPROM Programmable Logic Device.
- ◆ **Loader Version** – Version number of loader code.
- ◆ **Diagnostics Code Version** – Version of Power-On Self-Test (POST) and boot code.
- ◆ **Operation Code Version** – Version number of runtime code.
- ◆ **Thermal Detector** – The AS5700-54X has five detectors; the AS6700-32X has eight detectors.
- ◆ **Temperature** – Temperature at specified thermal detection point.

Web Interface

To view hardware and software version information.

1. Click System, then Switch.

Figure 4: General Switch Information



The screenshot shows a web interface for a switch configuration. At the top, it says "System > Switch". Below this, there are two sections: "Main Board Information" and "Management Software Information".

Main Board Information

Serial Number	571054X1452023
Number of Ports	54
Hardware Version	R01A
Main Power Status	Down
Redundant Power Status	Up

Management Software Information

Role	Master
EPLD Version	05/05/05
Loader Version	1.4.0.5
Diagnostics Code Version	
Operation Code Version	1.1.163.153

Temperature List Total: 5

Thermal Detector	Temperature (°C)
1	30
2	29
3	27
4	33
5	55

Configuring Support for Jumbo Frames

Use the System > Capability page to configure support for layer 2 jumbo frames. The switch provides more efficient throughput for large sequential data transfers by supporting jumbo frames of up to 9216 bytes for Gigabit, 10 Gigabit, and 40 Gigabit Ethernet ports or trunks. Compared to standard Ethernet frames that run only up to 1.5 KB, using jumbo frames significantly reduces the per-packet overhead required to process protocol encapsulation fields.

Usage Guidelines

- ◆ To use jumbo frames, both the source and destination end nodes (such as a computer or server) must support this feature. Also, when the connection is operating at full duplex, all switches in the network between the two end nodes must be able to accept the extended frame size. And for half-duplex connections, all devices in the collision domain would need to support jumbo frames.
- ◆ This command globally enables support for jumbo frames on all Gigabit and 10 Gigabit ports and trunks. To set the MTU for a specific interface, enable jumbo frames on this page, and then specify the required size of the MTU on the port or trunk interface configuration page (see ["Port Configuration" on page 104](#) or ["Trunk Configuration" on page 126](#)).

Parameters

The following parameters are displayed:

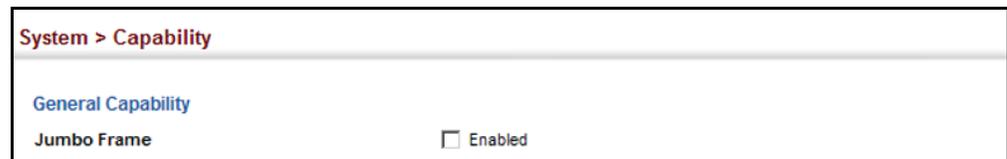
- ◆ **Jumbo Frame** – Configures support for jumbo frames. (Default: Disabled)

Web Interface

To configure support for jumbo frames:

1. Click System, then Capability.
2. Enable or disable support for jumbo frames.
3. Click Apply.

Figure 5: Configuring Support for Jumbo Frames



Displaying Bridge Extension Capabilities

Use the System > Capability page to display settings based on the Bridge MIB. The Bridge MIB includes extensions for managed devices that support Multicast Filtering, Traffic Classes, and Virtual LANs. You can access these extensions to display default settings for the key variables.

Parameters

The following parameters are displayed:

- ◆ **Extended Multicast Filtering Services** – This switch does not support the filtering of individual multicast addresses based on GMRP (GARP Multicast Registration Protocol).
- ◆ **Traffic Classes** – This switch provides mapping of user priorities to multiple traffic classes. (Refer to [“Class of Service” on page 199.](#))
- ◆ **Static Entry Individual Port** – This switch allows static filtering for unicast and multicast addresses. (Refer to [“Setting Static Addresses” on page 165.](#))
- ◆ **VLAN Version Number** – Based on IEEE 802.1Q, “1” indicates Bridges that support only single spanning tree (SST) operation, and “2” indicates Bridges that support multiple spanning tree (MST) operation.
- ◆ **VLAN Learning** – This switch uses Independent VLAN Learning (IVL), where each port maintains its own filtering database.
- ◆ **Local VLAN Capable** – This switch does not support multiple local bridges outside of the scope of 802.1Q defined VLANs.

- ◆ **Configurable PVID Tagging** – This switch allows you to override the default Port VLAN ID (PVID used in frame tags) and egress status (VLAN-Tagged or Untagged) on each port. (Refer to “VLAN Configuration” on page 145.)
- ◆ **Max Supported VLAN Numbers** – The maximum number of VLANs supported on this switch.
- ◆ **Max Supported VLAN ID** – The maximum configurable VLAN identifier supported on this switch.
- ◆ **GMRP** – GARP Multicast Registration Protocol (GMRP) allows network devices to register end stations with multicast groups. This switch does not support GMRP; it uses the Internet Group Management Protocol (IGMP) to provide automatic multicast filtering.

Web Interface

To view Bridge Extension information:

1. Click System, then Capability.

Figure 6: Displaying Bridge Extension Configuration



Managing System Files

This section describes how to upgrade the switch operating software or configuration files, and set the system start-up files.

Copying Files via FTP/TFTP or HTTP Use the System > File (Copy) page to upload/download firmware or configuration settings using FTP, TFTP or HTTP. By backing up a file to an FTP/TFTP server or management station, that file can later be downloaded to the switch to restore operation. Specify the method of file transfer, along with the file type and file names as required.

You can also set the switch to use new firmware or configuration settings without overwriting the current version. Just download the file using a different name from the current version, and then set the new file as the startup file.

Parameters

The following parameters are displayed:

- ◆ **Copy Type** – The firmware copy operation includes these options:
 - FTP Upload – Copies a file from an FTP server to the switch.
 - FTP Download – Copies a file from the switch to an FTP server.
 - HTTP Upload – Copies a file from a management station to the switch.
 - HTTP Download – Copies a file from the switch to a management station
 - TFTP Upload – Copies a file from a TFTP server to the switch.
 - TFTP Download – Copies a file from the switch to a TFTP server.
- ◆ **FTP/TFTP Server IP Address** – The IP address of an FTP/TFTP server.
- ◆ **User Name** – The user name for FTP server access.
- ◆ **Password** – The password for FTP server access.
- ◆ **File Type** – Specify Operation Code to copy firmware.
- ◆ **File Name** – The file name should not contain slashes (\ or /), the leading letter of the file name should not be a period (.), and the maximum length for file names is 32 characters for files on the switch or 127 characters for files on the server. (Valid characters: A-Z, a-z, 0-9, ":", "-", "_")



Note: Up to two copies of the system software (i.e., the runtime firmware) can be stored in the file directory on the switch.

Note: The maximum number of user-defined configuration files is limited only by available flash memory space.

Note: The file “Factory_Default_Config.cfg” can be copied to a file server or management station, but cannot be used as the destination file name on the switch.

Web Interface

To copy firmware files:

1. Click System, then File.
2. Select Copy from the Action list.
3. Select FTP Upload, HTTP Upload, or TFTP Upload as the file transfer method.
4. If FTP or TFTP Upload is used, enter the IP address of the file server.
5. If FTP Upload is used, enter the user name and password for your account on the FTP server.
6. Set the file type to Operation Code.
7. Enter the name of the file to upload.
8. Select a file on the switch to overwrite or specify a new file name.
9. Then click Apply.

Figure 7: Copy Firmware

The screenshot shows the 'System > File' web interface. At the top, there is a breadcrumb 'System > File'. Below it, the 'Action' dropdown is set to 'Copy'. The 'Copy Type' dropdown is set to 'TFTP Upload'. The 'TFTP Server IP Address' text box contains '192.168.1.99'. The 'File Type' dropdown is set to 'Operation Code'. The 'Source File Name' text box contains 'runtime.bix'. The 'Destination File Name' section has two radio buttons; the first is selected and has a dropdown menu showing 'runtime.bix'. The second radio button is unselected and has an empty text box next to it. At the bottom right, there are 'Apply' and 'Revert' buttons.

If you replaced a file currently used for startup and want to start using the new file, reboot the system via the System > Reset menu.

Saving the Running Configuration to a Local File

Use the System > File (Copy) page to save the current configuration settings to a local file on the switch. The configuration settings are not automatically saved by the system for subsequent use when the switch is rebooted. You must save these settings to the current startup file, or to another file which can be subsequently set as the startup file.

Parameters

The following parameters are displayed:

- ◆ **Copy Type** – The copy operation includes this option:
 - Running-Config – Copies the current configuration settings to a local file on the switch.
- ◆ **Destination File Name** – Copy to the currently designated startup file, or to a new file. The file name should not contain slashes (\ or /), the leading letter of the file name should not be a period (.), and the maximum length for file names is 32 characters for files on the switch. (Valid characters: A-Z, a-z, 0-9, “”, “-”, “_”)



Note: The maximum number of user-defined configuration files is limited only by available flash memory space.

Web Interface

To save the running configuration file:

1. Click System, then File.
2. Select Copy from the Action list.
3. Select Running-Config from the Copy Type list.
4. Select the current startup file on the switch to overwrite or specify a new file name.
5. Then click Apply.

Figure 8: Saving the Running Configuration

The screenshot shows the 'System > File' web interface. At the top, the breadcrumb 'System > File' is displayed. Below it, there is a form with the following elements:

- Action:** A dropdown menu with 'Copy' selected.
- Copy Type:** A dropdown menu with 'Running-Config' selected.
- Destination File Name:** A radio button is selected next to a dropdown menu showing 'startup1.cfg'. Below it is an empty text input field.
- Buttons:** 'Apply' and 'Revert' buttons are located at the bottom right of the form.

If you replaced a file currently used for startup and want to start using the new file, reboot the system via the System > Reset menu.

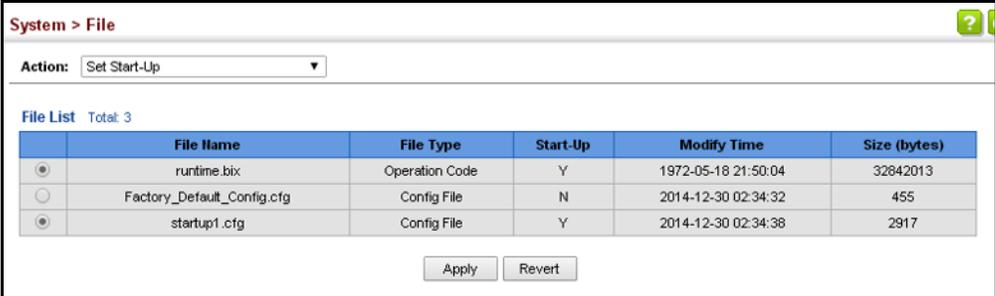
Setting The Start-Up File Use the System > File (Set Start-Up) page to specify the firmware or configuration file to use for system initialization.

Web Interface

To set a file to use for system initialization:

1. Click System, then File.
2. Select Set Start-Up from the Action list.
3. Mark the operation code or configuration file to be used at startup
4. Then click Apply.

Figure 9: Setting Start-Up Files



File Name	File Type	Start-Up	Modify Time	Size (bytes)
runtime.bix	Operation Code	Y	1972-05-18 21:50:04	32842013
Factory_Default_Config.cfg	Config File	N	2014-12-30 02:34:32	455
startup1.cfg	Config File	Y	2014-12-30 02:34:38	2917

To start using the new firmware or configuration settings, reboot the system via the System > Reset menu.

Showing System Files Use the System > File (Show) page to show the files in the system directory, or to delete a file.



Note: Files designated for start-up, and the Factory_Default_Config.cfg file, cannot be deleted.

Web Interface

To show the system files:

1. Click System, then File.
2. Select Show from the Action list.
3. To delete a file, mark it in the File List and click Delete.

Figure 10: Displaying System Files

System > File

Action: Show

File List Total: 3

<input type="checkbox"/>	File Name	File Type	Start-Up	Modify Time	Size (bytes)
<input type="checkbox"/>	runtime.bix	Operation Code	Y	1972-05-18 21:50:04	32842013
<input type="checkbox"/>	Factory_Default_Config.cfg	Config File	N	2014-12-30 02:34:32	455
<input type="checkbox"/>	startup1.cfg	Config File	Y	2014-12-30 02:34:38	2917

Delete Revert

Automatic Operation Code Upgrade

Use the System > File (Automatic Operation Code Upgrade) page to automatically download an operation code file when a file newer than the currently installed one is discovered on the file server. After the file is transferred from the server and successfully written to the file system, it is automatically set as the startup file, and the switch is rebooted.

Usage Guidelines

- ◆ If this feature is enabled, the switch searches the defined URL once during the bootup sequence.
- ◆ FTP (port 21) and TFTP (port 69) are both supported. Note that the TCP/UDP port bindings cannot be modified to support servers listening on non-standard ports.
- ◆ The host portion of the upgrade file location URL must be a valid IPv4 IP address. DNS host names are not recognized. Valid IP addresses consist of four numbers, 0 to 255, separated by periods.
- ◆ The path to the directory must also be defined. If the file is stored in the root directory for the FTP/TFTP service, then use the "/" to indicate this (e.g., ftp://192.168.0.1/).
- ◆ The file name must not be included in the upgrade file location URL. The file name of the code stored on the remote server must be aos5700-54x.bix (using lower case letters exactly as indicated here). Enter the file name for other switches described in this manual exactly as shown on the web interface.
- ◆ The FTP connection is made with PASV mode enabled. PASV mode is needed to traverse some fire walls, even if FTP traffic is not blocked. PASV mode cannot be disabled.
- ◆ The switch-based search function is case-insensitive in that it will accept a file name in upper or lower case (i.e., the switch will accept AOS5700-54X.BIX from the server even though AOS5700-54X.bix was requested). However, keep in mind that the file systems of many operating systems such as Unix and most Unix-like systems (FreeBSD, NetBSD, OpenBSD, and most Linux distributions, etc.) are case-sensitive, meaning that two files in the same directory,

aos5700-54x.bix and *AOS5700-54X.BIX* are considered to be unique files. Thus, if the upgrade file is stored as *AOS5700-54X.BIX* (or even *Aos5700-54x.bix*) on a case-sensitive server, then the switch (requesting *AOS5700-54X.bix*) will not be upgraded because the server does not recognize the requested file name and the stored file name as being equal. A notable exception in the list of case-sensitive Unix-like operating systems is Mac OS X, which by default is case-insensitive. Please check the documentation for your server's operating system if you are unsure of its file system's behavior.

- ◆ Note that the switch itself does not distinguish between upper and lower-case file names, and only checks to see if the file stored on the server is more recent than the current runtime image.
- ◆ If two operation code image files are already stored on the switch's file system, then the non-startup image is deleted before the upgrade image is transferred.
- ◆ The automatic upgrade process will take place in the background without impeding normal operations (data switching, etc.) of the switch.
- ◆ During the automatic search and transfer process, the administrator cannot transfer or update another operation code image, configuration file, public key, or HTTPS certificate (i.e., no other concurrent file management operations are possible).
- ◆ The upgrade operation code image is set as the startup image after it has been successfully written to the file system.
- ◆ The switch will send an SNMP trap and make a log entry upon all upgrade successes and failures.
- ◆ The switch will immediately restart after the upgrade file is successfully written to the file system and set as the startup image.

Parameters

The following parameters are displayed:

- ◆ **Automatic Opcode Upgrade** – Enables the switch to search for an upgraded operation code file during the switch bootup process. (Default: Disabled)
- ◆ **Automatic Upgrade Location URL** – Defines where the switch should search for the operation code upgrade file. The last character of this URL must be a forward slash ("/"). The *ECS4660-28F.bix* filename must not be included since it is automatically appended by the switch. (Options: ftp, tftp)

The following syntax must be observed:

tftp://host[/filedir]

- **tftp://** – Defines TFTP protocol for the server connection.

- *host* – Defines the IP address of the TFTP server. Valid IP addresses consist of four numbers, 0 to 255, separated by periods. DNS host names are not recognized.
- *filedir* – Defines the directory, relative to the TFTP server root, where the upgrade file can be found. Nested directory structures are accepted. The directory name must be separated from the host, and in nested directory structures, from the parent directory, with a prepended forward slash “/”.
- / – The forward slash must be the last character of the URL.

ftp://[*username[:password@]*]*host*[/*filedir*]/

- **ftp://** – Defines FTP protocol for the server connection.
- *username* – Defines the user name for the FTP connection. If the user name is omitted, then “anonymous” is the assumed user name for the connection.
- *password* – Defines the password for the FTP connection. To differentiate the password from the user name and host portions of the URL, a colon (:) must precede the password, and an “at” symbol (@), must follow the password. If the password is omitted, then “” (an empty string) is the assumed password for the connection.
- *host* – Defines the IP address of the FTP server. Valid IP addresses consist of four numbers, 0 to 255, separated by periods. DNS host names are not recognized.
- *filedir* – Defines the directory, relative to the FTP server root, where the upgrade file can be found. Nested directory structures are accepted. The directory name must be separated from the host, and in nested directory structures, from the parent directory, with a prepended forward slash “/”.
- / – The forward slash must be the last character of the URL.

Examples

The following examples demonstrate the URL syntax for a TFTP server at IP address 192.168.0.1 with the operation code image stored in various locations:

- `tftp://192.168.0.1/`
The image file is in the TFTP root directory.
- `tftp://192.168.0.1/switch-opcode/`
The image file is in the “switch-opcode” directory, relative to the TFTP root.
- `tftp://192.168.0.1/switches/opcode/`
The image file is in the “opcode” directory, which is within the “switches” parent directory, relative to the TFTP root.

The following examples demonstrate the URL syntax for an FTP server at IP address 192.168.0.1 with various user name, password and file location options presented:

- ftp://192.168.0.1/
The user name and password are empty, so “anonymous” will be the user name and the password will be blank. The image file is in the FTP root directory.
- ftp://switches:upgrade@192.168.0.1/
The user name is “switches” and the password is “upgrade”. The image file is in the FTP root.
- ftp://switches:upgrade@192.168.0.1/switches/opcode/
The user name is “switches” and the password is “upgrade”. The image file is in the “opcode” directory, which is within the “switches” parent directory, relative to the FTP root.

Web Interface

To configure automatic code upgrade:

1. Click System, then File.
2. Select Automatic Operation Code Upgrade from the Action list.
3. Mark the check box to enable Automatic Opcode Upgrade.
4. Enter the URL of the FTP or TFTP server, and the path and directory containing the operation code.
5. Click Apply.

Figure 11: Configuring Automatic Code Upgrade

The screenshot shows a web interface for configuring automatic code upgrade. The breadcrumb path is "System > File". The "Action" dropdown menu is set to "Automatic Operation Code Upgrade". Below this, there is a section for "Automatic Opcode Upgrade" with a checkbox labeled "Enabled" that is currently unchecked. The "Automatic Upgrade Location URL" field contains the text "ftp://admin:admin@192.168.0.1/opfiles/". A note below the URL field states: "Note: For automatic upgrades, the operation code file name must be set as aos5700-54x.bix." At the bottom right of the form, there are two buttons: "Apply" and "Revert".

If a new image is found at the specified location, the following type of messages will be displayed during bootup.

```
.  
. .  
Automatic Upgrade is looking for a new image  
New image detected: current version 1.1.1.0; new version 1.1.1.2  
Image upgrade in progress  
The switch will restart after upgrade succeeds  
Downloading new image  
Flash programming started  
Flash programming completed  
The switch will now restart  
. . .
```

Setting the System Clock

Simple Network Time Protocol (SNTP) allows the switch to set its internal clock based on periodic updates from a time server (SNTP or NTP). Maintaining an accurate time on the switch enables the system log to record meaningful dates and times for event entries. You can also manually set the clock. If the clock is not set manually or via SNTP, the switch will only record the time from the factory default set at the last bootup.

When the SNTP client is enabled, the switch periodically sends a request for a time update to a configured time server. You can configure up to three time server IP addresses. The switch will attempt to poll each server in the configured sequence.

Setting the Time Manually Use the System > Time (Configure General - Manual) page to set the system time on the switch manually without using SNTP.

Parameters

The following parameters are displayed:

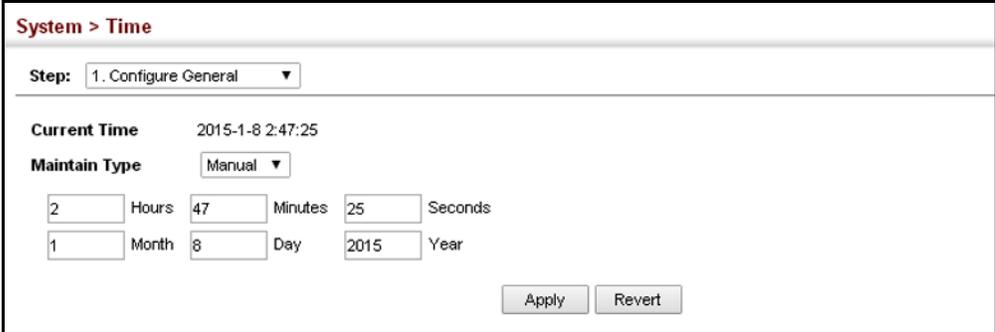
- ◆ **Current Time** – Shows the current time set on the switch.
- ◆ **Hours** – Sets the hour. (Range: 0-23)
- ◆ **Minutes** – Sets the minute value. (Range: 0-59)
- ◆ **Seconds** – Sets the second value. (Range: 0-59)
- ◆ **Month** – Sets the month. (Range: 1-12)
- ◆ **Day** – Sets the day of the month. (Range: 1-31)
- ◆ **Year** – Sets the year. (Range: 1970-2037)

Web Interface

To manually set the system clock:

1. Click System, then Time.
2. Select Configure General from the Step list.
3. Select Manual from the Maintain Type list.
4. Enter the time and date in the appropriate fields.
5. Click Apply

Figure 12: Manually Setting the System Clock



The screenshot shows a web interface titled "System > Time". At the top, there is a "Step:" dropdown menu set to "1. Configure General". Below this, the "Current Time" is displayed as "2015-1-8 2:47:25". The "Maintain Type" is set to "Manual" via a dropdown menu. There are two rows of input fields: the first row contains "2" for Hours, "47" for Minutes, and "25" for Seconds; the second row contains "1" for Month, "8" for Day, and "2015" for Year. At the bottom right, there are two buttons: "Apply" and "Revert".

Setting the SNTP Polling Interval Use the System > Time (Configure General - SNTP) page to set the polling interval at which the switch will query the specified time servers.

Parameters

The following parameters are displayed:

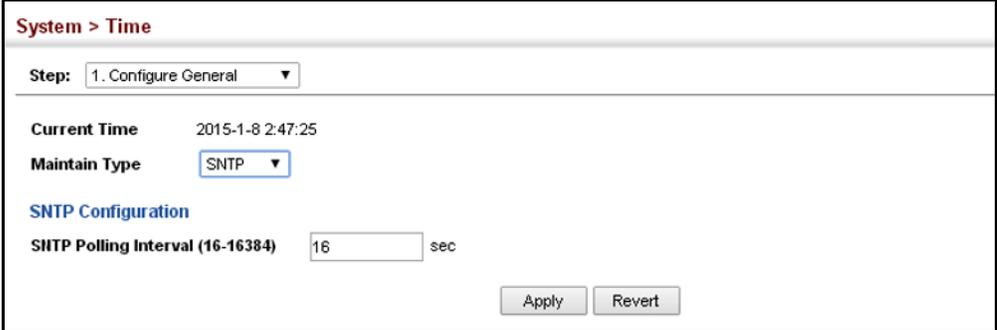
- ◆ **Current Time** – Shows the current time set on the switch.
- ◆ **SNTP Polling Interval** – Sets the interval between sending requests for a time update from a time server. (Range: 16-16384 seconds; Default: 16 seconds)

Web Interface

To set the polling interval for SNTP:

1. Click System, then Time.
2. Select Configure General from the Step list.
3. Select SNTP from the Maintain Type list.
4. Modify the polling interval if required.
5. Click Apply

Figure 13: Setting the Polling Interval for SNTP



The screenshot shows a web interface for configuring system time. At the top, it says "System > Time". Below that, there is a "Step:" dropdown menu set to "1. Configure General". The "Current Time" is displayed as "2015-1-8 2:47:25". The "Maintain Type" is set to "SNTP" in a dropdown menu. Under the "SNTP Configuration" section, the "SNTP Polling Interval (16-16384)" is set to "16" seconds. At the bottom right, there are "Apply" and "Revert" buttons.

Configuring NTP Use the System > Time (Configure General - NTP) page to configure NTP authentication and show the polling interval at which the switch will query the specified time servers.

Parameters

The following parameters are displayed:

- ◆ **Current Time** – Shows the current time set on the switch.
- ◆ **Authentication Status** – Enables authentication for time requests and updates between the switch and NTP servers. (Default: Disabled)

You can enable NTP authentication to ensure that reliable updates are received from only authorized NTP servers. The authentication keys and their associated key number must be centrally managed and manually distributed to NTP servers and clients. The key numbers and key values must match on both the server and client.
- ◆ **Polling Interval** – Shows the interval between sending requests for a time update from NTP servers. (Fixed: 1024 seconds)

Web Interface

To set the clock maintenance type to NTP:

1. Click System, then Time.
2. Select Configure General from the Step list.
3. Select NTP from the Maintain Type list.
4. Enable authentication if required.
5. Click Apply

Figure 14: Configuring NTP

The screenshot shows a web interface for configuring NTP. At the top, it says "System > Time". Below that is a "Step:" dropdown menu set to "1. Configure General". The "Current Time" is displayed as "2015-1-8 2:47:25". The "Maintain Type" is set to "NTP" via a dropdown menu. Under the "NTP Configuration" section, there is an "Authentication Status" checkbox labeled "Enabled" which is currently unchecked. The "Polling Interval" is set to "1024 sec". At the bottom right, there are "Apply" and "Revert" buttons.

Configuring Time Servers Use the System > Time (Configure Time Server) pages to specify the IP address for NTP/SNTP time servers, or to set the authentication key for NTP time servers.

Specifying SNTP Time Servers

Use the System > Time (Configure Time Server – Configure SNTP Server) page to specify the IP address for up to three SNTP time servers.

Parameters

The following parameters are displayed:

- ◆ **SNTP Server IP Address** – Sets the IPv4 or IPv6 address for up to three time servers. The switch attempts to update the time from the first server, if this fails it attempts an update from the next server in the sequence.

Web Interface

To set the SNTP time servers:

1. Click System, then Time.
2. Select Configure Time Server from the Step list.
3. Select Configure SNTP Server from the Action list.
4. Enter the IP address of up to three time servers.
5. Click Apply.

Figure 15: Specifying SNTP Time Servers

The screenshot shows the 'System > Time' configuration page. At the top, there is a breadcrumb 'System > Time'. Below it, there are two dropdown menus: 'Step: 2. Configure Time Server' and 'Action: Configure SNTP Server'. The main area contains three input fields for SNTP Server IP addresses: 'SNTP Server IP Address 1' with the value '10.1.0.19', 'SNTP Server IP Address 2' with the value '137.62.140.80', and 'SNTP Server IP Address 3' with the value '128.250.36.2'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

Specifying NTP Time Servers

Use the System > Time (Configure Time Server – Add NTP Server) page to add the IP address for up to 50 NTP time servers.

Parameters

The following parameters are displayed:

- ◆ **NTP Server IP Address** – Adds the IPv4 or IPv6 address for up to 50 time servers. The switch will poll the specified time servers for updates when the clock maintenance type is set to NTP on the System > Time (Configure General) page. It issues time synchronization requests at a fixed interval of 1024 seconds. The switch will poll all the time servers configured, the responses received are filtered and compared to determine the most reliable and accurate time update for the switch.
- ◆ **Version** – Specifies the NTP version supported by the server. (Fixed: Version 3)
- ◆ **Authentication Key** – Specifies the number of the key in the NTP Authentication Key List to use for authentication with the configured server. NTP authentication is optional. If enabled on the System > Time (Configure General) page, you must also configure at least one key on the System > Time (Add NTP Authentication Key) page. (Range: 1-65535)

Web Interface

To add an NTP time server to the server list:

1. Click System, then Time.
2. Select Configure Time Server from the Step list.
3. Select Add NTP Server from the Action list.
4. Enter the IP address of an NTP time server, and specify the index of the authentication key if authentication is required.
5. Click Apply.

Figure 16: Adding an NTP Time Server

The screenshot shows the 'System > Time' configuration page. At the top, there is a breadcrumb 'System > Time'. Below it, there are two dropdown menus: 'Step: 2. Configure Time Server' and 'Action: Add NTP Server'. The main form contains three input fields: 'NTP Server IP Address' with the value '192.168.3.20', 'Version' with the value '3', and 'Authentication Key (1-65535)' with the value '3' and '(optional)' text to its right. At the bottom right of the form are two buttons: 'Apply' and 'Revert'.

To show the list of configured NTP time servers:

1. Click System, then Time.
2. Select Configure Time Server from the Step list.
3. Select Show NTP Server from the Action list.

Figure 17: Showing the NTP Time Server List

The screenshot shows the 'System > Time' configuration page with the 'Action' dropdown set to 'Show NTP Server'. Below the form, there is a table titled 'NTP Server List' with a 'Total: 1' indicator. The table has four columns: a checkbox, 'Server IP Address', 'Version', and 'Authentication Key'. There is one row of data with the values: checkbox (unchecked), '192.168.3.20', '3', and '3'. Below the table are two buttons: 'Delete' and 'Revert'.

<input type="checkbox"/>	Server IP Address	Version	Authentication Key
<input type="checkbox"/>	192.168.3.20	3	3

Specifying NTP Authentication Keys

Use the System > Time (Configure Time Server – Add NTP Authentication Key) page to add an entry to the authentication key list.

Parameters

The following parameters are displayed:

- ◆ **Authentication Key** – Specifies the number of the key in the NTP Authentication Key List to use for authentication with a configured server. NTP authentication is optional. When enabled on the System > Time (Configure General) page, you must also configure at least one key on this page. Up to 255 keys can be configured on the switch. (Range: 1-65535)
- ◆ **Key Context** – An MD5 authentication key string. The key string can be up to 32 case-sensitive printable ASCII characters (no spaces).
NTP authentication key numbers and values must match on both the server and client.

Web Interface

To add an entry to NTP authentication key list:

1. Click System, then Time.
2. Select Configure Time Server from the Step list.
3. Select Add NTP Authentication Key from the Action list.
4. Enter the index number and MD5 authentication key string.
5. Click Apply.

Figure 18: Adding an NTP Authentication Key

To show the list of configured NTP authentication keys:

1. Click System, then Time.
2. Select Configure Time Server from the Step list.
3. Select Show NTP Authentication Key from the Action list.

Figure 19: Showing the NTP Authentication Key List

NTP Authentication Key List		Total: 1
<input type="checkbox"/>	Authentication Key	Key Context
<input type="checkbox"/>	3	8J0774Q6699747D10867F12S505J62770084708278G1357878N8475052113Q69137L8

Setting the Time Zone Use the System > Time (Configure Time Server) page to set the time zone. SNTP uses Coordinated Universal Time (or UTC, formerly Greenwich Mean Time, or GMT) based on the time at the Earth's prime meridian, zero degrees longitude, which passes through Greenwich, England. To display a time corresponding to your local time, you must indicate the number of hours and minutes your time zone is east (before) or west (after) of UTC. You can choose one of the 80 predefined time zone definitions, or you can manually configure the parameters for your local time zone.

Parameters

The following parameters are displayed:

- ◆ **Predefined Configuration** – A drop-down box provides access to the 80 predefined time zone configurations. Each choice indicates its offset from UTC and lists at least one major city or location covered by the time zone.
- ◆ **User-defined Configuration** – Allows the user to define all parameters of the local time zone.
 - **Direction** – Configures the time zone to be before (east of) or after (west of) UTC.
 - **Name** – Assigns a name to the time zone. (Range: 1-30 characters)
 - **Hours** (0-13) – The number of hours before or after UTC. The maximum value before UTC is 12. The maximum value after UTC is 13.
 - **Minutes** (0-59) – The number of minutes before/after UTC.

Web Interface

To set your local time zone:

1. Click System, then Time.
2. Select Configure Time Zone from the Step list.
3. Set the offset for your time zone relative to the UTC in hours and minutes.
4. Click Apply.

Figure 20: Setting the Time Zone

The screenshot shows a web interface for configuring the system time zone. The breadcrumb is 'System > Time'. The current step is '3. Configure Time Zone'. There are two radio buttons: 'Predefined Configuration' (unselected) and 'User Defined Configuration' (selected). Under 'User Defined Configuration', there are four fields: 'Direction' (a dropdown menu set to 'After UTC'), 'Name' (a text input field containing 'UTC'), 'Hours (0-13)' (a text input field containing '0'), and 'Minutes (0-59)' (a text input field containing '0'). Below these fields is a note: 'Note: The maximum value before UTC is 12:00. The maximum value after UTC is 13:00.' At the bottom right, there are two buttons: 'Apply' and 'Revert'.

Configuring The Console Port

Use the System > Console menu to configure connection parameters for the switch's console port. You can access the onboard configuration program by attaching a VT100 compatible device to the switch's serial console port. Management access through the console port is controlled by various parameters, including a password (only configurable through the CLI), time outs, and basic communication settings. Note that these parameters can be configured via the web or CLI interface.

Parameters

The following parameters are displayed:

- ◆ **Login Timeout** – Sets the interval that the system waits for a user to log into the CLI. If a login attempt is not detected within the timeout interval, the connection is terminated for the session. (Range: 10-300 seconds; Default: 300 seconds)
- ◆ **Exec Timeout** – Sets the interval that the system waits until user input is detected. If user input is not detected within the timeout interval, the current session is terminated. (Range: 60-65535 seconds; Default: 600 seconds)
- ◆ **Password Threshold** – Sets the password intrusion threshold, which limits the number of failed logon attempts. When the logon attempt threshold is reached, the system interface becomes silent for a specified amount of time (set by the Silent Time parameter) before allowing the next logon attempt. (Range: 1-120; Default: 3 attempts)
- ◆ **Silent Time** – Sets the amount of time the management console is inaccessible after the number of unsuccessful logon attempts has been exceeded. (Range: 1-65535 seconds; Default: Disabled)

- ◆ **Data Bits** – Sets the number of data bits per character that are interpreted and generated by the console port. If parity is being generated, specify 7 data bits per character. If no parity is required, specify 8 data bits per character. (Range: 1-8; Default: 8 bits)
- ◆ **Stop Bits** – Sets the number of the stop bits transmitted per byte. (Range: 1-2; Default: 1 stop bit)
- ◆ **Parity** – Defines the generation of a parity bit. Communication protocols provided by some terminals can require a specific parity bit setting. Specify Even, Odd, or None. (Default: None)
- ◆ **Speed** – Sets the terminal line's baud rate for transmit (to terminal) and receive (from terminal). Set the speed to match the baud rate of the device connected to the serial port. (Range: 9600, 19200, 38400, 57600, or 115200 baud; Default: 115200 baud)



Note: The password for the console connection can only be configured through the CLI (see “password” in the *CLI Reference Guide*).

Note: Password checking can be enabled or disabled for logging in to the console connection (see “login” in the *CLI Reference Guide*). You can select authentication by a single global password as configured for the password command, or by passwords set up for specific user-name accounts. The default is for local passwords configured on the switch.

Web Interface

To configure parameters for the console port:

1. Click System, then Console.
2. Specify the connection parameters as required.
3. Click Apply

Figure 21: Console Port Settings

The screenshot shows the 'System > Console' configuration page. It includes the following settings:

- Login Timeout (10-300):** 300 sec
- Exec Timeout (60-65535):** 600 sec
- Password Threshold (1-120):** 3
- Silent Time (1-65535):** sec
- Data Bits:** 8
- Stop Bits:** 1
- Parity:** None
- Speed:** 115200 baud

Buttons for 'Apply' and 'Revert' are located at the bottom right of the configuration area.

Configuring Telnet Settings

Use the System > Telnet menu to configure parameters for accessing the CLI over a Telnet connection. You can access the onboard configuration program over the network using Telnet (i.e., a virtual terminal). Management access via Telnet can be enabled/disabled and other parameters set, including the TCP port number, time outs, and a password. Note that the password is only configurable through the CLI.) These parameters can be configured via the web or CLI interface.

Parameters

The following parameters are displayed:

- ◆ **Telnet Status** – Enables or disables Telnet access to the switch. (Default: Enabled)
- ◆ **TCP Port** – Sets the TCP port number for Telnet on the switch. (Range: 1-65535; Default: 23)
- ◆ **Max Sessions** – Sets the maximum number of Telnet sessions that can simultaneously connect to this system. (Range: 0-8; Default: 8)
A maximum of eight sessions can be concurrently opened for Telnet and Secure Shell (i.e., both Telnet and SSH share a maximum number of eight sessions).
- ◆ **Login Timeout** – Sets the interval that the system waits for a user to log into the CLI. If a login attempt is not detected within the timeout interval, the connection is terminated for the session. (Range: 10-300 seconds; Default: 300 seconds)
- ◆ **Exec Timeout** – Sets the interval that the system waits until user input is detected. If user input is not detected within the timeout interval, the current session is terminated. (Range: 60-65535 seconds; Default: 600 seconds)

- ◆ **Password Threshold** – Sets the password intrusion threshold, which limits the number of failed logon attempts. When the logon attempt threshold is reached, the system interface becomes silent for a specified amount of time (set by the Silent Time parameter) before allowing the next logon attempt. (Range: 1-120; Default: 3 attempts)
- ◆ **Silent Time** – Sets the amount of time the management interface is inaccessible after the number of unsuccessful logon attempts has been exceeded. (Range: 1-65535 seconds; Default: Disabled)



Note: The password for the console connection can only be configured through the CLI (see “password” in the *CLI Reference Guide*).

Note: Password checking can be enabled or disabled for login to the console connection (see the “login” command in the *CLI Reference Guide*). You can select authentication by a single global password as configured for the password command, or by passwords set up for specific user-name accounts. The default is for local passwords configured on the switch.

Web Interface

To configure parameters for the console port:

1. Click System, then Telnet.
2. Specify the connection parameters as required.
3. Click Apply

Figure 22: Telnet Connection Settings

System > Telnet	
Telnet Status	<input checked="" type="checkbox"/> Enabled
TCP Port (1-65535)	<input type="text" value="23"/>
Max Sessions (0-8)	<input type="text" value="8"/>
Login Timeout (10-300)	<input type="text" value="300"/> sec
Exec Timeout (60-65535)	<input type="text" value="600"/> sec
Password Threshold (1-120)	<input checked="" type="checkbox"/> <input type="text" value="3"/>
Silent Time (1-65535)	<input type="checkbox"/> <input type="text"/> sec

Displaying CPU Utilization

Use the System > CPU Utilization page to display information on CPU utilization.

Parameters

The following parameters are displayed:

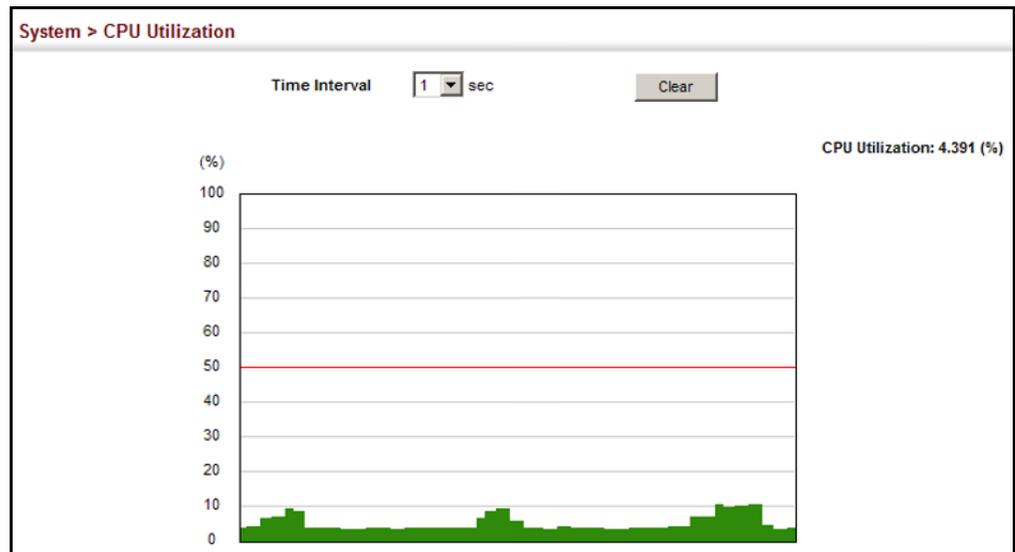
- ◆ **Time Interval** – The interval at which to update the displayed utilization rate. (Options: 1, 5, 10, 30, 60 seconds; Default: 1 second)
- ◆ **CPU Utilization** – CPU utilization over specified interval.

Web Interface

To display CPU utilization:

1. Click System, then CPU Utilization.
2. Change the update interval if required. Note that the interval is changed as soon as a new setting is selected.

Figure 23: Displaying CPU Utilization



Displaying Memory Utilization

Use the System > Memory Status page to display memory utilization parameters.

Parameters

The following parameters are displayed:

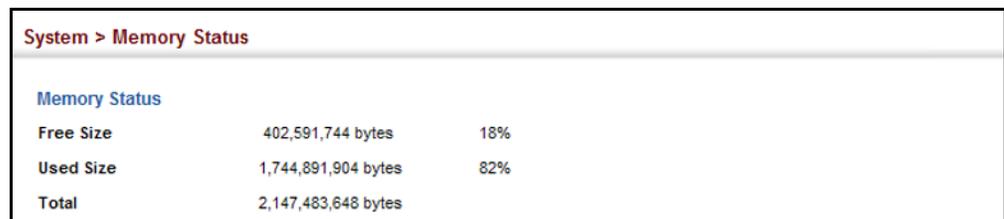
- ◆ **Free Size** – The amount of memory currently free for use.
- ◆ **Used Size** – The amount of memory allocated to active processes.
- ◆ **Total** – The total amount of system memory.

Web Interface

To display memory utilization:

1. Click System, then Memory Status.

Figure 24: Displaying Memory Utilization



System > Memory Status		
Memory Status		
Free Size	402,591,744 bytes	18%
Used Size	1,744,891,904 bytes	82%
Total	2,147,483,648 bytes	

Resetting the System

Use the System > Reset menu to restart the switch immediately, at a specified time, after a specified delay, or at a periodic interval.

Command Usage

- ◆ This command resets the entire system.
- ◆ When the system is restarted, it will always run the Power-On Self-Test. It will also retain all configuration information stored in non-volatile memory by the “copy running-config startup-config” command in the *CLI Reference Guide*.

Parameters

System Reload Information

The following parameters are displayed by clicking Show in the Action menu:

- ◆ **Reload Information** – Displays information on the next scheduled reload and selected reload mode.

- ◆ **Delete** – Deletes the marked entry.
- ◆ **Revert** – Cancels the current settings.

System Reload Configuration

The following parameters are displayed by clicking Configure in the Action menu:

- ◆ **Reset Mode** – Restarts the switch immediately or at the specified time(s).
 - **Immediately** – Restarts the system immediately.
 - **In** – Specifies an interval after which to reload the switch. (The specified time must be equal to or less than 24 days.)
 - *hours* – The number of hours, combined with the minutes, before the switch resets. (Range: 0-576)
 - *minutes* – The number of minutes, combined with the hours, before the switch resets. (Range: 0-59)
 - **At** – Specifies a time at which to reload the switch.
 - DD - The day of the month at which to reload. (Range: 01-31)
 - MM - The month at which to reload. (Range: 01-12)
 - YYYY - The year at which to reload. (Range: 1970-2037)
 - HH - The hour at which to reload. (Range: 00-23)
 - MM - The minute at which to reload. (Range: 00-59)
 - **Regularly** – Specifies a periodic interval at which to reload the switch.

Time

- HH - The hour at which to reload. (Range: 00-23)
- MM - The minute at which to reload. (Range: 00-59)

Period

- Daily - Every day.
 - Weekly - Day of the week at which to reload. (Range: Sunday ... Saturday)
 - Monthly - Day of the month at which to reload. (Range: 1-31)
- ◆ **Apply** – Click this button to save the current configuration settings.

- ◆ **Revert** – Click this button to cancel any changes.

Web Interface

To restart the switch:

1. Click System, then Reset.
2. Select the required reset mode.
3. For any option other than to reset immediately, fill in the required parameters
4. Click Apply.
5. When prompted, confirm that you want reset the switch.

Figure 25: Restarting the Switch (Immediately)

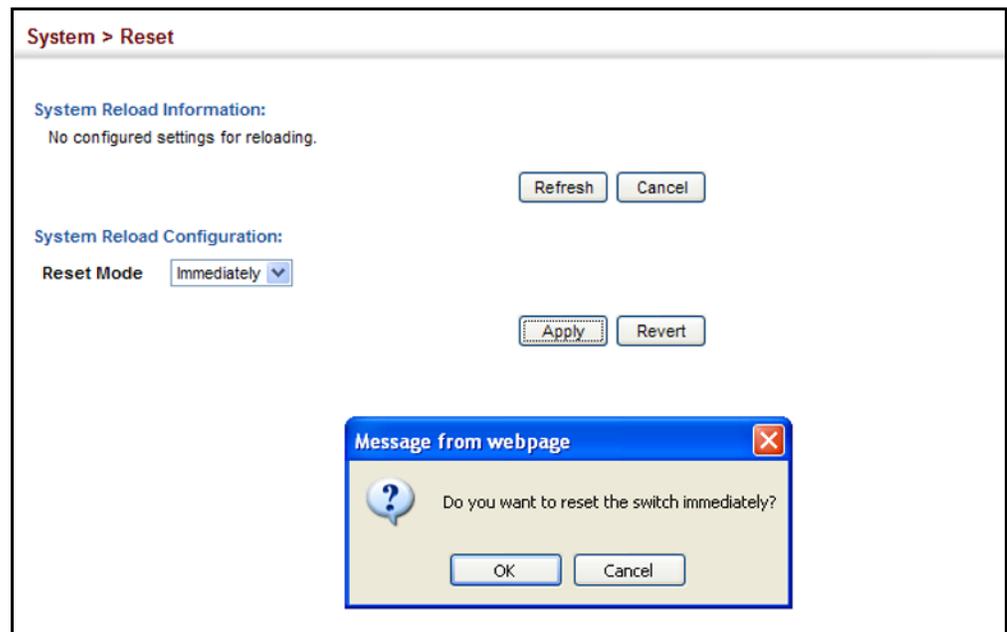


Figure 26: Restarting the Switch (In)

System > Reset

System Reload Information:
The switch will be rebooted at Jan 1 02:54:25 2001. Remaining Time: 0 days, 1 hours, 10 minutes, 0 seconds.
Reloading switch in time: 5 hours 26 minutes.
Reloading switch regularity time: 11:20 everyday.

System Reload Configuration:
Reset Mode
Reload switch in hours minutes.
Note: The specified time must be equal to or less than 24 days.

Figure 27: Restarting the Switch (At)

System > Reset

System Reload Information:
The switch will be rebooted at Jan 1 02:54:25 2001. Remaining Time: 0 days, 1 hours, 10 minutes, 0 seconds.
Reloading switch in time: 5 hours 26 minutes.
Reloading switch regularity time: 11:20 everyday.

System Reload Configuration:
Reset Mode
Reload switch at (DD/MM/YYYY) (HH:MM)
Warning: You have to setup system time first. Otherwise this function won't work.

Figure 28: Restarting the Switch (Regularly)

System > Reset

System Reload Information:
No configured settings for reloading.

System Reload Configuration:
Reset Mode
Time (HH:MM)
Period
 Daily
 Weekly
 Monthly

Warning: You have to setup system time first. Otherwise this function won't work.

4

Interface Configuration

This chapter describes the following topics:

- ◆ [Port Configuration](#) – Configures connection settings, including auto-negotiation, or manual setting of speed, duplex mode, and flow control.
- ◆ [Local Port Mirroring](#) – Sets the source and target ports for mirroring on the local switch.
- ◆ [Remote Port Mirroring](#) – Configures mirroring of traffic from remote switches for analysis at a destination port on the local switch.
- ◆ [Displaying Statistics](#) – Shows Interface, Etherlike, and RMON port statistics in table or chart form.
- ◆ [Displaying Statistical History](#) – Displays statistical history for the specified interfaces.
- ◆ [Displaying Transceiver Data](#) – Displays identifying information, and operational parameters for optical transceivers which support DDM.
- ◆ [Configuring Transceiver Thresholds](#) – Configures thresholds for alarm and warning messages for optical transceivers which support DDM.
- ◆ [Trunk Configuration](#) – Configures static or dynamic trunks.
- ◆ [Traffic Segmentation](#) – Configures the uplinks and down links to a segmented group of ports.

Port Configuration

This section describes how to configure port connections, mirror traffic from one port to another, and run cable diagnostics.

Configuring by Port List Use the Interface > Port > General (Configure by Port List) page to enable/disable an interface, set auto-negotiation and the interface capabilities to advertise, or manually fix the speed, duplex mode, and flow control.

Command Usage

- ◆ 10GBASE-SFP+ connections are fixed at 10G - full duplex, and 40GBASE-QSFP+ connections at 40G - full duplex. Auto-negotiation must be disabled before you can configure or force an RJ-45 interface to use the Flow Control option.
- ◆ When using auto-negotiation¹, the optimal settings will be negotiated between the link partners based on their advertised capabilities. To set flow control and symmetric pause frames under auto-negotiation, the required operation modes must be specified in the capabilities list for an interface.
- ◆ The Speed/Duplex mode is fixed at 100full for 100BASE-FX transceivers, 1000full for Gigabit transceivers, and 10Gfull for 10 Gigabit transceivers. When auto-negotiation is enabled¹, the only attributes which can be advertised include flow control and symmetric pause frames.

Using Jumbo Frames

- ◆ Use the jumbo frame attribute on the System > Capability page to enable or disable jumbo frames for all 10 Gigabit and 40 Gigabit Ethernet ports. Then specify the required MTU size for a specific interface on the port configuration page.
- ◆ The comparison of packet size against the configured port MTU considers only the incoming packet size, and is not affected by the fact that an ingress port is a tagged port or a QinQ ingress port. In other words, any additional size (for example, a tagged field of 4 bytes added by the chip) will not be considered when comparing the egress packet's size against the configured MTU.
- ◆ When pinging the switch from an external device, information added for the Ethernet header can increase the packet size by at least 42 bytes for an untagged packet, and 46 bytes for a tagged packet. If the adjusted frame size exceeds the configured port MTU, the switch will not respond to the ping message.
- ◆ For other traffic types, calculation of overall frame size is basically the same, including the additional header fields SA(6) + DA(6) + Type(2) + VLAN-Tag(4) (for tagged packets, for untagged packets, the 4-byte field will not be added by

1. Support for auto-negotiation depends on transceiver type, such as 1G SFP.

switch), and the payload. This should all be less than the configured port MTU, including the CRC at the end of the frame.

- ◆ For QinQ, the overall frame size is still calculated as described above, and does not add the length of the second tag to the frame.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-32/54)
- ◆ **Type** – Indicates the port type. (1000BASE SFP, 10GBASE SFP+, 40GBASE QSFP)
- ◆ **Name** – Allows you to label an interface. (Range: 1-64 characters)
- ◆ **Admin** – Allows you to manually disable an interface. You can disable an interface due to abnormal behavior (e.g., excessive collisions), and then re-enable it after the problem has been resolved. You may also disable an interface for security reasons. (Default: Enabled)
- ◆ **Media Type** – Configures the forced transceiver mode for SFP+ ports.
 - **None** - Forced transceiver mode is not used for SFP+ ports. (This is the default setting for RJ-45 ports and SFP+ ports.)
 - **SFP-Forced 1000SFP** - Always uses the SFP+ port at 1000 Mbps, full duplex.
 - **SFP-Forced 10GSFP** - Always uses the SFP+ port at 10 Gbps, full duplex.
- ◆ **Autonegotiation** (Port Capabilities) – Not supported on this switch. Forced mode is used for all ports.
 Default: Autonegotiation disabled;
 Forced mode capabilities for -
 1000BASE-SX/LX (SFP+) – 1000full
 10GBASE-CR/SR/LR/LRM (SFP+) – 10Gfull
 40GBASE-T-CR4 (QSFP+) – 40Gfull
- ◆ **Speed/Duplex** – Shows the port speed and duplex mode.
- ◆ **Flow Control** – Allows automatic or manual selection of flow control.
- ◆ **MTU Size** – The maximum transfer unit (MTU) allowed for layer 2 packets crossing a 1G/10G/40G Ethernet port or trunk. (Range: 1500-12288 bytes; Default: 1518 bytes)
- ◆ **Link Up Down Trap** – Issues a notification message whenever a port link is established or broken. (Default: Enabled)

Web Interface

To configure port connection parameters:

1. Click Interface, Port, General.
2. Select Configure by Port List from the Action List.
3. Modify the required interface settings.
4. Click Apply.

Figure 29: Configuring Connections by Port List

Port	Type	Name	Admin	Media Type	Autonegotiation	Speed Duplex	Flow Control	MTU Size (1500-12288)	Link Up Down Trap
1	10GBASE SFP+		<input checked="" type="checkbox"/> Enabled	None	<input type="checkbox"/> Enabled <input checked="" type="checkbox"/> 10Gf <input type="checkbox"/> FC <input type="checkbox"/> Sym	10Gfull	<input checked="" type="checkbox"/> Enabled	1518	<input checked="" type="checkbox"/> Enabled
2	10GBASE SFP+		<input checked="" type="checkbox"/> Enabled	None	<input type="checkbox"/> Enabled <input checked="" type="checkbox"/> 10Gf <input type="checkbox"/> FC <input type="checkbox"/> Sym	10Gfull	<input type="checkbox"/> Enabled	1518	<input checked="" type="checkbox"/> Enabled
3	10GBASE SFP+		<input checked="" type="checkbox"/> Enabled	None	<input type="checkbox"/> Enabled <input checked="" type="checkbox"/> 10Gf <input type="checkbox"/> FC <input type="checkbox"/> Sym	10Gfull	<input type="checkbox"/> Enabled	1518	<input checked="" type="checkbox"/> Enabled

Configuring by Port Range Use the Interface > Port > General (Configure by Port Range) page to enable/disable an interface, set auto-negotiation and the interface capabilities to advertise, or manually fix the speed, duplex mode, and flow control.

For more information on command usage and a description of the parameters, refer to [“Configuring by Port List” on page 104](#).

Web Interface

To configure port connection parameters:

1. Click Interface, Port, General.
2. Select Configure by Port Range from the Action List.
3. Enter to range of ports to which your configuration changes apply.
4. Modify the required interface settings.
5. Click Apply.

Figure 30: Configuring Connections by Port Range

Displaying Connection Status Use the Interface > Port > General (Show Information) page to display the current connection status, including link state, speed/duplex mode, flow control, and auto-negotiation.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier.
- ◆ **Type** – Indicates the port type. (1000BASE SFP, 10GBASE SFP+, 40GBASE QSFP)
- ◆ **Name** – Interface label.
- ◆ **Admin** – Shows if the port is enabled or disabled.
- ◆ **Oper Status** – Indicates if the link is Up or Down.
- ◆ **Media Type** – Shows the forced transceiver mode.
- ◆ **Autonegotiation** – Shows that auto-negotiation is disabled.
- ◆ **Oper Speed Duplex** – Shows the current speed and duplex mode.
- ◆ **Oper Flow Control** – Shows the flow control type used.
- ◆ **MTU Size** – The maximum transfer unit (MTU) allowed for layer 2 packets crossing a Gigabit or 10 Gigabit Ethernet port or trunk.
- ◆ **Link Up/Down Trap** – Shows if link-up or link-down notifications are enabled.

Web Interface

To display port connection parameters:

1. Click Interface, Port, General.
2. Select Show Information from the Action List.

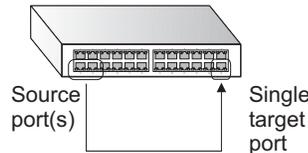
Figure 31: Displaying Port Information

Port	Type	Name	Admin	Oper Status	Media Type	Autonegotiation	Oper Speed Duplex	Oper Flow Control	MTU Size	Link Up Down Trap
1	10GBASE SFP+		Enabled	Down	None	Disabled	10Gfull	None	1518	Enabled
2	10GBASE SFP+		Enabled	Down	None	Disabled	10Gfull	None	1518	Enabled
3	10GBASE SFP+		Enabled	Down	None	Disabled	10Gfull	None	1518	Enabled

Configuring Local Port Mirroring

Use the Interface > Port > Mirror page to mirror traffic from any source port to a target port for real-time analysis. You can then attach a logic analyzer or RMON probe to the target port and study the traffic crossing the source port in a completely unobtrusive manner.

Figure 32: Configuring Local Port Mirroring



Command Usage

- ◆ Traffic can be mirrored from one or more source ports to a destination port on the same switch (local port mirroring as described in this section).
- ◆ Monitor port speed should match or exceed source port speed, otherwise traffic may be dropped from the monitor port.
- ◆ When mirroring port traffic, the target port must be included in the same VLAN as the source port when using MSTP (see “Spanning Tree Algorithm” on page 173).
- ◆ Note that Spanning Tree BPDU packets are not mirrored to the target port.

Parameters

These parameters are displayed:

- ◆ **Source Port** – The port whose traffic will be monitored.

- ◆ **Target Port** – The port that will mirror the traffic on the source port.
- ◆ **Type** – Allows you to select which traffic to mirror to the target port, Rx (receive), Tx (transmit), or Both. (Default: Both)

Web Interface

To configure a local mirror session:

1. Click Interface, Port, Mirror.
2. Select Add from the Action List.
3. Specify the source port.
4. Specify the monitor port.
5. Specify the traffic type to be mirrored.
6. Click Apply.

Figure 33: Configuring Local Port Mirroring

Interface > Port > Mirror

Action: Add

Source Port Unit 1 Port 7

Target Port Unit 1 Port 8

Type Rx

Apply Revert

To display the configured mirror sessions:

1. Click Interface, Port, Mirror.
2. Select Show from the Action List.

Figure 34: Displaying Local Port Mirror Sessions

Interface > Port > Mirror

Action: Show

Mirror Session List Total: 3

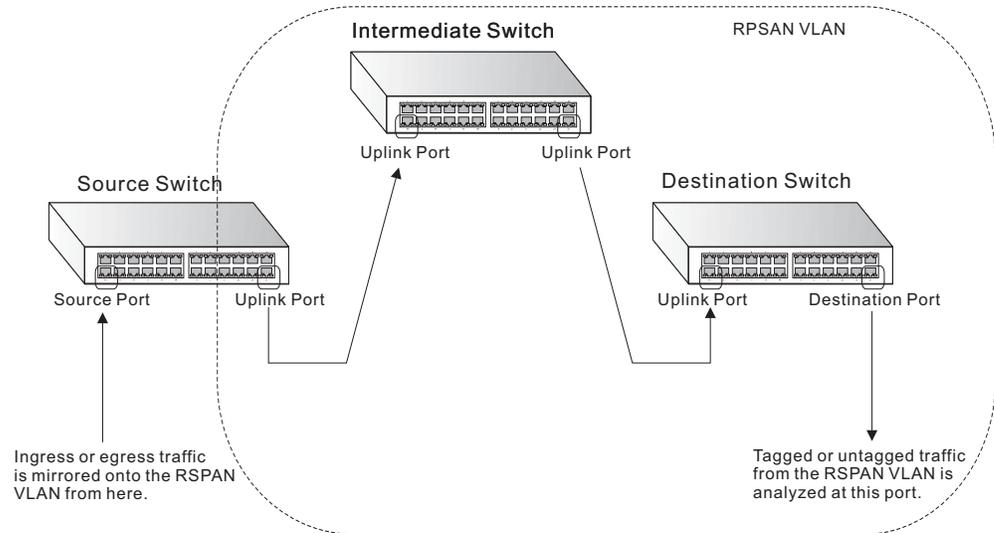
	Source (Unit/Port)	Target (Unit/Port)	Type
<input type="checkbox"/>	1 / 1	1 / 5	Rx
<input type="checkbox"/>	1 / 2	1 / 5	Rx
<input type="checkbox"/>	1 / 3	1 / 5	Rx

Delete Revert

Configuring Remote Port Mirroring

Use the Interface > Port > RSPAN page to mirror traffic from remote switches for analysis at a destination port on the local switch. This feature, also called Remote Switched Port Analyzer (RSPAN), carries traffic generated on the specified source ports for each session over a user-specified VLAN dedicated to that RSPAN session in all participating switches. Monitored traffic from one or more sources is copied onto the RSPAN VLAN through IEEE 802.1Q trunk or hybrid ports that carry it to any RSPAN destination port monitoring the RSPAN VLAN as shown in the figure below.

Figure 35: Configuring Remote Port Mirroring



Command Usage

- ◆ Traffic can be mirrored from one or more source ports to a destination port on the same switch (local port mirroring as described in [“Configuring Local Port Mirroring” on page 108](#)), or from one or more source ports on remote switches to a destination port on this switch (remote port mirroring as described in this section).
- ◆ *Configuration Guidelines*

Take the following step to configure an RSPAN session:

1. Use the VLAN Static List (see [“Configuring VLAN Groups” on page 147](#)) to reserve a VLAN for use by RSPAN (marking the “Remote VLAN” field on this page. (Default VLAN 1 is prohibited.)
2. Set up the source switch on the RSPAN configuration page by specifying the mirror session, the switch’s role (Source), the RSPAN VLAN, and the uplink port. Then specify the source port(s), and the traffic type to monitor (Rx, Tx or Both).
3. Set up all intermediate switches on the RSPAN configuration page, entering the mirror session, the switch’s role (Intermediate), the RSPAN VLAN, and the uplink port(s).

4. Set up the destination switch on the RSPAN configuration page by specifying the mirror session, the switch's role (Destination), the destination port², whether or not the traffic exiting this port will be tagged or untagged, and the RSPAN VLAN. Then specify each uplink port where the mirrored traffic is being received.

◆ *RSPAN Limitations*

The following limitations apply to the use of RSPAN on this switch:

- *RSPAN Ports* – Only ports can be configured as an RSPAN source, destination, or uplink; static and dynamic trunks are not allowed. A port can only be configured as one type of RSPAN interface – source, destination, or uplink. Also, note that the source port and destination port cannot be configured on the same switch.
- *Local/Remote Mirror* – The destination of a local mirror session (created on the Interface > Port > Mirror page) cannot be used as the destination for RSPAN traffic.
- *Spanning Tree* – If the spanning tree is disabled, BPDUs will not be flooded onto the RSPAN VLAN.
- *MAC address learning* is not supported on RSPAN uplink ports when RSPAN is enabled on the switch. Therefore, even if spanning tree is enabled after RSPAN has been configured, MAC address learning will still not be re-started on the RSPAN uplink ports.
- *IEEE 802.1X* – RSPAN and 802.1X are mutually exclusive functions. When 802.1X is enabled globally, RSPAN uplink ports cannot be configured, even though RSPAN source and destination ports can still be configured. When RSPAN uplink ports are enabled on the switch, 802.1X cannot be enabled globally.
- *Port Security* – If port security is enabled on any port, that port cannot be set as an RSPAN uplink port, even though it can still be configured as an RSPAN source or destination port. Also, when a port is configured as an RSPAN uplink port, port security cannot be enabled on that port.

Parameters

These parameters are displayed:

- ◆ **Session** – A number identifying this RSPAN session. (Range: 1-2)

Only two mirror sessions are allowed, including both local and remote mirroring. If local mirroring is enabled (see [page 108](#)), then there is only one session available for RSPAN.

2. If the RSPAN packets are configured as untagged, then the destination port type cannot be trunk mode (see [“Adding Static Members to VLANs” on page 150](#)).

- ◆ **Operation Status** – Indicates whether or not RSPAN is currently functioning.
- ◆ **Switch Role** – Specifies the role this switch performs in mirroring traffic.
 - **None** – This switch will not participate in RSPAN.
 - **Source** - Specifies this device as the source of remotely mirrored traffic.
 - **Intermediate** - Specifies this device as an intermediate switch, transparently passing mirrored traffic from one or more sources to one or more destinations.
 - **Destination** - Specifies this device as a switch configured with a destination port which is to receive mirrored traffic for this session.
- ◆ **Remote VLAN** – The VLAN to which traffic mirrored from the source port will be flooded. The VLAN specified in this field must first be reserved for the RSPAN application using the VLAN > Static page (see [page 150](#)).
- ◆ **Uplink Port** – A port on any switch participating in RSPAN through which mirrored traffic is passed on to or received from the RSPAN VLAN.

Only one uplink port can be configured on a source switch, but there is no limitation on the number of uplink ports configured on an intermediate or destination switch.

Only destination and uplink ports will be assigned by the switch as members of the RSPAN VLAN. Ports cannot be manually assigned to an RSPAN VLAN through the VLAN > Static page. Nor can GVRP dynamically add port members to an RSPAN VLAN. Also, note that the VLAN > Static (Show) page will not display any members for an RSPAN VLAN, but will only show configured RSPAN VLAN identifiers.
- ◆ **Source Port** – Specifies the ports to monitor. (Switch Role as Source)
- ◆ **Type** – Specifies the traffic type to be mirrored remotely. (Options: Rx, Tx, Both)
- ◆ **Port (Uplink)** – Specifies the destination port to monitor the traffic mirrored from the source ports. Only one destination port can be configured on the same switch per session, but a destination port can be configured on more than one switch for the same session. Also note that a destination port can still send and receive switched traffic, and participate in any Layer 2 protocols to which it has been assigned. (Switch Role as Intermediate or Destination)
- ◆ **Tag** – Specifies whether or not the traffic exiting the destination port to the monitoring device carries the RSPAN VLAN tag.

Web Interface

To configure a remote mirror session:

1. Click Interface, RSPAN.
2. Set the Switch Role to None, Source, Intermediate, or Destination.

3. Configure the required settings for each switch participating in the RSPAN VLAN.
4. Click Apply.

Figure 36: Configuring Remote Port Mirroring (Source)

Interface > RSPAN

Session: 1

Operation Status: Up

Switch Role: Source

Remote VLAN: 2

Uplink Port: 1

Source Port Configuration List Total: 52

Source Port	Type
1	None
2	Both
3	None
4	None
5	None

Figure 37: Configuring Remote Port Mirroring (Intermediate)

Interface > RSPAN

Session: 1

Operation Status: Up

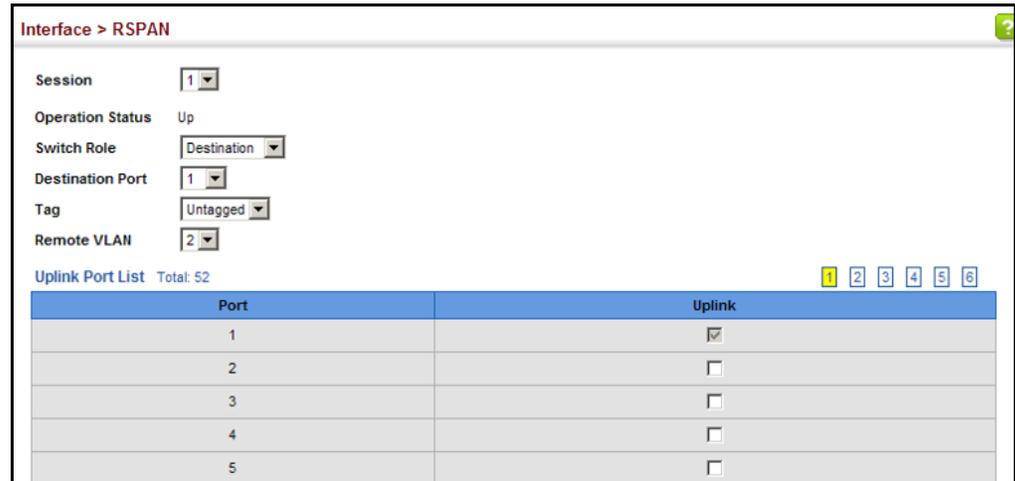
Switch Role: Intermediate

Remote VLAN: 2

Uplink Port List Total: 52

Port	Uplink
1	<input checked="" type="checkbox"/>
2	<input type="checkbox"/>
3	<input type="checkbox"/>
4	<input type="checkbox"/>
5	<input type="checkbox"/>

Figure 38: Configuring Remote Port Mirroring (Destination)



Showing Port or Trunk Statistics

Use the Interface > Port/Trunk > Statistics or Chart page to display standard statistics on network traffic from the Interfaces Group and Ethernet-like MIBs, as well as a detailed breakdown of traffic based on the RMON MIB. Interfaces and Ethernet-like statistics display errors on the traffic passing through each port. This information can be used to identify potential problems with the switch (such as a faulty port or unusually heavy traffic). RMON statistics provide access to a broad range of statistics, including a total count of different frame types and sizes passing through each port. All values displayed have been accumulated since the last system reboot, and are shown as counts per second. Statistics are refreshed every 60 seconds by default.



Note: RMON groups 2, 3 and 9 can only be accessed using SNMP management software.

Parameters

These parameters are displayed:

Table 6: Port Statistics

Parameter	Description
<i>Interface Statistics</i>	
Received Octets	The total number of octets received on the interface, including framing characters.
Transmitted Octets	The total number of octets transmitted out of the interface, including framing characters.
Received Errors	The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
Transmitted Errors	The number of outbound packets that could not be transmitted because of errors.

Table 6: Port Statistics (Continued)

Parameter	Description
Received Unicast Packets	The number of subnetwork-unicast packets delivered to a higher-layer protocol.
Transmitted Unicast Packets	The total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent.
Received Discarded Packets	The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space.
Transmitted Discarded Packets	The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space.
Received Multicast Packets	The number of packets, delivered by this sub-layer to a higher (sub-)layer, which were addressed to a multicast address at this sub-layer.
Transmitted Multicast Packets	The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sub-layer, including those that were discarded or not sent.
Received Broadcast Packets	The number of packets, delivered by this sub-layer to a higher (sub-)layer, which were addressed to a broadcast address at this sub-layer.
Transmitted Broadcast Packets	The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a broadcast address at this sub-layer, including those that were discarded or not sent.
Received Unknown Packets	The number of packets received via the interface which were discarded because of an unknown or unsupported protocol.
<i>Etherlike Statistics</i>	
Single Collision Frames	The number of successfully transmitted frames for which transmission is inhibited by exactly one collision.
Multiple Collision Frames	A count of successfully transmitted frames for which transmission is inhibited by more than one collision.
Late Collisions	The number of times that a collision is detected later than 512 bit-times into the transmission of a packet.
Excessive Collisions	A count of frames for which transmission on a particular interface fails due to excessive collisions. This counter does not increment when the interface is operating in full-duplex mode.
Deferred Transmissions	A count of frames for which the first transmission attempt on a particular interface is delayed because the medium was busy.
Frames Too Long	A count of frames received on a particular interface that exceed the maximum permitted frame size.
Alignment Errors	The number of alignment errors (missynchronized data packets).
FCS Errors	A count of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check. This count does not include frames received with frame-too-long or frame-too-short error.
SQE Test Errors	A count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular interface.
Carrier Sense Errors	The number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame.

Table 6: Port Statistics (Continued)

Parameter	Description
Internal MAC Receive Errors	A count of frames for which reception on a particular interface fails due to an internal MAC sublayer receive error.
Internal MAC Transmit Errors	A count of frames for which transmission on a particular interface fails due to an internal MAC sublayer transmit error.
<i>RMON Statistics</i>	
Drop Events	The total number of events in which packets were dropped due to lack of resources.
Jabbers	The total number of frames received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS or alignment error.
Fragments	The total number of frames received that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS or alignment error.
Collisions	The best estimate of the total number of collisions on this Ethernet segment.
Received Octets	Total number of octets of data received on the network. This statistic can be used as a reasonable indication of Ethernet utilization.
Received Packets	The total number of packets (bad, broadcast and multicast) received.
Broadcast Packets	The total number of good packets received that were directed to the broadcast address. Note that this does not include multicast packets.
Multicast Packets	The total number of good packets received that were directed to this multicast address.
Undersize Packets	The total number of packets received that were less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.
Oversize Packets	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.
64 Bytes Packets	The total number of packets (including bad packets) received and transmitted that were 64 octets in length (excluding framing bits but including FCS octets).
65-127 Byte Packets	The total number of packets (including bad packets) received and transmitted where the number of octets fall within the specified range (excluding framing bits but including FCS octets).
128-255 Byte Packets	
256-511 Byte Packets	
512-1023 Byte Packets	
1024-1518 Byte Packets	
1519-1536 Byte Packets	
<i>Utilization Statistics</i>	
Input Octets in kbits per second	Number of octets entering this interface in kbits/second.
Input Packets per second	Number of packets entering this interface per second.
Input Utilization	The input utilization rate for this interface.
Output Octets in kbits per second	Number of octets leaving this interface in kbits/second.

Table 6: Port Statistics (Continued)

Parameter	Description
Output Packets per second	Number of packets leaving this interface per second.
Output Utilization	The output utilization rate for this interface.

Web Interface

To show a list of port statistics:

1. Click Interface, Port, Statistics.
2. Select the statistics mode to display (Interface, Etherlike, RMON or Utilization).
3. Select a port from the drop-down list.
4. Use the Refresh button at the bottom of the page if you need to update the screen.

Figure 39: Showing Port Statistics (Table)

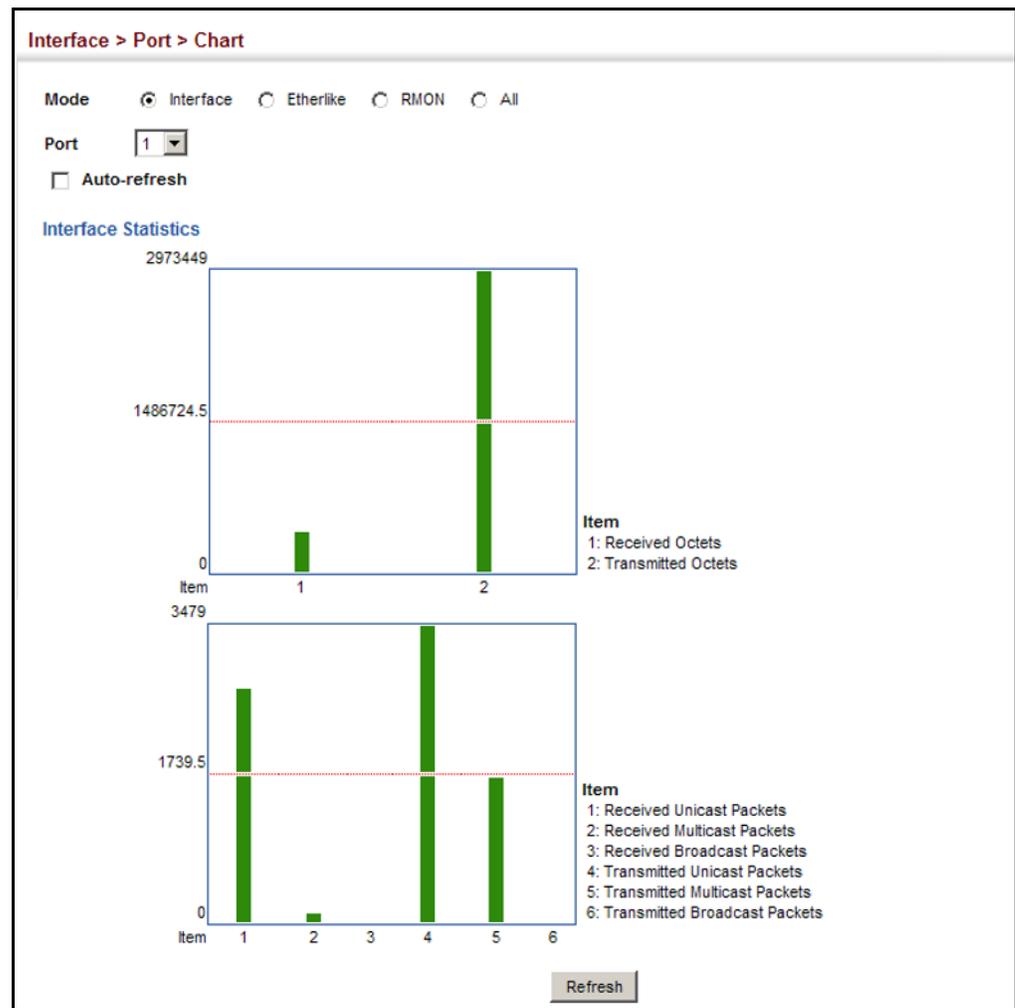
The screenshot shows the web interface for 'Interface > Port > Statistics'. At the top, there are radio buttons for 'Mode' with 'Interface' selected, and 'Etherlike', 'RMON', and 'Utilization' unselected. Below that is a 'Port' dropdown menu showing '1' and an 'Auto-refresh' checkbox which is unchecked. The main content is a table titled 'Interface Statistics' with two columns of data. At the bottom right, there is a 'Refresh' button.

Interface Statistics			
Received Octets	182057	Transmitted Octets	1353652
Received Errors	0	Transmitted Errors	0
Received Unicast Packets	1270	Transmitted Unicast Packets	1700
Received Discarded Packets	0	Transmitted Discarded Packets	0
Received Multicast Packets	9	Transmitted Multicast Packets	838
Received Broadcast Packets	23	Transmitted Broadcast Packets	2
Received Unknown Packets	0		

To show a chart of port statistics:

1. Click Interface, Port, Chart.
2. Select the statistics mode to display (Interface, Etherlike, RMON or All).
3. If Interface, Etherlike, RMON statistics mode is chosen, select a port from the drop-down list. If All (ports) statistics mode is chosen, select the statistics type to display.

Figure 40: Showing Port Statistics (Chart)



Displaying Statistical History Use the Interface > Port > History or Interface > Trunk > History page to display statistical history for the specified interfaces.

Command Usage

For a description of the statistics displayed on these pages, see [“Showing Port or Trunk Statistics” on page 114](#).

Parameters

These parameters are displayed:

Add

- ◆ **Port** – Port number. (Range: 1-32/54)
- ◆ **History Name** – Name of sample interval. (Range: 1-32 characters)
- ◆ **Interval** - The interval for sampling statistics. (Range: 1-86400 minutes)
- ◆ **Requested Buckets** - The number of samples to take. (Range: 1-96)

Show

- ◆ **Port** – Port number. (Range: 1-32/54)
- ◆ **History Name** – Name of sample interval. (Default settings: 15min, 1day)
- ◆ **Interval** - The interval for sampling statistics.
- ◆ **Requested Buckets** - The number of samples to take.

Show Details

- ◆ **Mode**
 - **Status** – Shows the sample parameters.
 - **Current Entry** – Shows current statistics for the specified port and named sample.
 - **Input Previous Entries** – Shows statistical history for ingress traffic.
 - **Output Previous Entries** – Shows statistical history for egress traffic.
- ◆ **Port** – Port number. (Range: 1-32/54)
- ◆ **Name** – Name of sample interval.

Web Interface

To configure a periodic sample of statistics:

1. Click Interface, Port, Statistics, or Interface, Trunk, Statistics.
2. Select Add from the Action menu.
3. Select an interface from the Port or Trunk list.
4. Enter the sample name, the interval, and the number of buckets requested.
5. Click Apply.

Figure 41: Configuring a History Sample

Interface > Port > History

Action: Add

Port: 1

History Name: rd#1

Interval (1-86400): 60

Requested Buckets (1-96): 50

Apply Revert

To show the configured entries for a history sample:

1. Click Interface, Port, Statistics, or Interface, Trunk, Statistics.
2. Select Show from the Action menu.
3. Select an interface from the Port or Trunk list.

Figure 42: Showing Entries for History Sampling

Interface > Port > History

Action: Show

Port: 1

History Name List Total: 3

<input type="checkbox"/>	History Name	Interval	Requested Buckets
<input type="checkbox"/>	15min	900	96
<input type="checkbox"/>	1day	86400	7
<input type="checkbox"/>	rd#1	60	50

Delete Revert

To show the configured parameters for a sampling entry:

1. Click Interface, Port, Statistics, or Interface, Trunk, Statistics.
2. Select Show Details from the Action menu.
3. Select Status from the options for Mode.
4. Select an interface from the Port or Trunk list.
5. Select an sampling entry from the Name list.

Figure 43: Showing Status of Statistical History Sample

Interface > Port > History

Action: Show Details

Mode Status Current Entry Input Previous Entries Output Previous Entries

Port 1

Name 15min

History Status

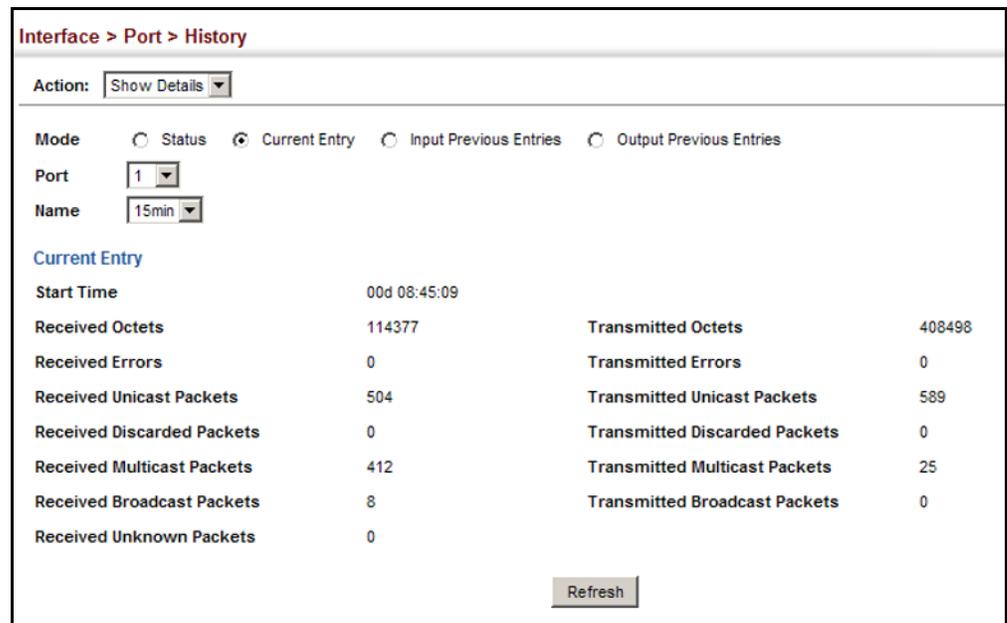
Name	15min
Interval	15 minute(s)
Requested Buckets	96
Granted Buckets	35
Status	Active

Refresh

To show statistics for the current interval of a sample entry:

1. Click Interface, Port, Statistics, or Interface, Trunk, Statistics.
2. Select Show Details from the Action menu.
3. Select Current Entry from the options for Mode.
4. Select an interface from the Port or Trunk list.
5. Select an sampling entry from the Name list.

Figure 44: Showing Current Statistics for a History Sample



To show ingress or egress traffic statistics for a sample entry:

1. Click Interface, Port, Statistics, or Interface, Trunk, Statistics.
2. Select Show Details from the Action menu.
3. Select Input Previous Entry or Output Previous Entry from the options for Mode.
4. Select an interface from the Port or Trunk list.
5. Select an sampling entry from the Name list.

Figure 45: Showing Ingress Statistics for a History Sample

Interface > Port > History

Action: Show Details ▾

Mode: Status Current Entry Input Previous Entries Output Previous Entries

Port: 1 ▾

Time: 15min ▾

Input Previous Entry List Total: 26

Start Time	%	Octets	Unicast	Multicast	Broadcast	Discarded	Errors
00d 00:00:00	0.00	50136	6	485	240	0	0
00d 00:15:01	0.00	45047	0	481	214	0	0
00d 00:30:01	0.00	155934	954	481	229	0	0
00d 00:45:01	0.00	128467	662	481	217	0	0
00d 01:00:01	0.00	130588	671	481	221	0	0
00d 01:15:01	0.00	81077	278	481	229	0	0
00d 01:30:01	0.00	135199	774	481	236	0	0
00d 01:45:01	0.00	155762	872	481	213	0	0
00d 02:00:01	0.00	128586	651	480	225	0	0
00d 02:15:01	0.00	127251	646	481	214	0	0

Refresh

Displaying Transceiver Data

Use the Interface > Port > Transceiver page to display identifying information, and operational for optical transceivers which support Digital Diagnostic Monitoring (DDM).

Parameters

These parameters are displayed:

- ◆ **Port** – Port number. (Range: 1-32/54)
- ◆ **General** – Information on connector type and vendor-related parameters.
- ◆ **DDM Information** – Information on temperature, supply voltage, laser bias current, laser power, and received optical power.

The switch can display diagnostic information for SFP modules which support the SFF-8472 Specification for Diagnostic Monitoring Interface for Optical Transceivers. This information allows administrators to remotely diagnose problems with optical devices. This feature, referred to as Digital Diagnostic Monitoring (DDM) provides information on transceiver parameters.

Web Interface

To display identifying information and functional parameters for optical transceivers:

1. Click Interface, Port, Transceiver.
2. Select a port from the scroll-down list.

Figure 46: Displaying Transceiver Data

The screenshot shows a web interface for configuring a transceiver. At the top, it says 'Interface > Port > Transceiver'. Below this, there is a dropdown menu for 'Port' with '11' selected. The interface is divided into two main sections: 'General' and 'DDM Information'. The 'General' section lists various parameters such as Connector Type (LC), Fiber Type (Multimode 50um (M5), Multimode 62.5um (M6)), Baud Rate (2100 MBd), Vendor OUI (00-90-65), Vendor Name (FINISAR CORP.), Vendor PN (FTLF8519P3BTL), Vendor Rev (A), Vendor SN (PKM1XUU), and Date Code (11-05-25). The 'DDM Information' section lists Temperature (1.50 °C), Vcc (0.00 V), Bias Current (0.00 mA (ch1), 0.00 mA (ch2), 0.00 mA (ch3), 0.00 mA (ch4)), and RX Power (-40.00 dBm (ch1), -40.00 dBm (ch2), -40.00 dBm (ch3), -40.00 dBm (ch4)).

Interface > Port > Transceiver	
Port	11
General	
Connector Type	LC
Fiber Type	Multimode 50um (M5), Multimode 62.5um (M6)
Eth Compliance Codes	1000BASE-SX
Baud Rate	2100 MBd
Vendor OUI	00-90-65
Vendor Name	FINISAR CORP.
Vendor PN	FTLF8519P3BTL
Vendor Rev	A
Vendor SN	PKM1XUU
Date Code	11-05-25
DDM Information	
Temperature	1.50 °C
Vcc	0.00 V
Bias Current	0.00 mA (ch1), 0.00 mA (ch2), 0.00 mA (ch3), 0.00 mA (ch4)
RX Power	-40.00 dBm (ch1), -40.00 dBm (ch2), -40.00 dBm (ch3), -40.00 dBm (ch4)

Configuring Transceiver Thresholds Use the Interface > Port > Transceiver page to configure thresholds for alarm and warning messages for optical transceivers which support DDM.

Parameters

These parameters are displayed:

- ◆ **Port** – Port number. (Range: 1-32/54)
- ◆ **General** – Information on connector type and vendor-related parameters.
- ◆ **DDM Information** – Information on temperature, supply voltage, laser bias current, laser power, and received optical power.

The switch can display diagnostic information for SFP modules which support the SFF-8472 Specification for Diagnostic Monitoring Interface for Optical Transceivers. This information allows administrators to remotely diagnose problems with optical devices. This feature, referred to as Digital Diagnostic Monitoring (DDM) provides information on transceiver parameters.
- ◆ **Trap** – Sends a trap when any of the transceiver’s operation values falls outside of specified thresholds. (Default: Disabled)
- ◆ **Auto Mode** – Uses default threshold settings obtained from the transceiver to determine when an alarm or trap message should be sent. (Default: Enabled)

- ◆ **DDM Thresholds** – Information on alarm and warning thresholds. The switch can be configured to send a trap when the measured parameter falls outside of the specified thresholds.

The following alarm and warning parameters are supported:

- **High Alarm** – Sends an alarm message when the high threshold is crossed.
- **High Warning** – Sends a warning message when the high threshold is crossed.
- **Low Warning** – Sends a warning message when the low threshold is crossed.
- **Low Alarm** – Sends an alarm message when the low threshold is crossed.

The configurable ranges are:

- **Temperature:** -200.00-200.00 °C
- **Voltage:** 1.00-255.00 Volts
- **Current:** 1.00-255.00 mA
- **Power:** -99.99-99.99 dBm

The threshold value for Rx and Tx power is calculated as the power ratio in decibels (dB) of the measured power referenced to one milliwatt (mW).

Threshold values for alarm and warning messages can be configured as described below.

- A high-threshold alarm or warning message is sent if the current value is greater than or equal to the threshold, and the last sample value was less than the threshold. After a rising event has been generated, another such event will not be generated until the sampled value has fallen below the high threshold and reaches the low threshold.
- A low-threshold alarm or warning message is sent if the current value is less than or equal to the threshold, and the last sample value was greater than the threshold. After a falling event has been generated, another such event will not be generated until the sampled value has risen above the low threshold and reaches the high threshold.
- Threshold events are triggered as described above to avoid a hysteresis effect which would continuously trigger event messages, for example, if the power level were to fluctuate just above and below either the high threshold or the low threshold.
- Trap messages configured by this command are sent to any management station configured as an SNMP trap manager using the Administration > SNMP (Configure Trap) page.

Web Interface

To configure threshold values for optical transceivers:

1. Click Interface, Port, Transceiver.
2. Select a port from the scroll-down list.
3. Set the switch to send a trap based on default or manual settings.
4. Set alarm and warning thresholds if manual configuration is used.
5. Click Apply.

Figure 47: Configuring Transceiver Thresholds

	Low Alarm	Low Warning	High Warning	High Alarm
Temperature(°C)	35.03	0.00	0.00	13.00
Voltage(Volts)	0.20	2.88	1.67	0.00
Current(mA)	5.15	36.01	0.00	33.34
Rx Power(dBm)	-0.85	-0.85	-0.85	1.50

Restore Default Click this button to restore default DDM thresholds values.

Apply Revert

Trunk Configuration

This section describes how to configure static and dynamic trunks.

You can create multiple links between devices that work as one virtual, aggregate link. A trunk offers a dramatic increase in bandwidth for network segments where bottlenecks exist, as well as providing a fault-tolerant link between two devices. You can create up to 27/16 trunks on the AOS5700-54X and AOS6700-32X at the same time on the switch.

The switch supports both static trunking and dynamic Link Aggregation Control Protocol (LACP). Static trunks have to be manually configured at both ends of the link, and the switches must comply with the Cisco EtherChannel standard. On the other hand, LACP configured ports can automatically negotiate a trunked link with LACP-configured ports on another device. You can configure any number of ports on the switch as LACP, as long as they are not already configured as part of a static trunk. If ports on another device are also configured as LACP, the switch and the other device will negotiate a trunk link between them. If an LACP trunk consists of more than eight ports, all other ports will be placed in standby mode. Should one

link in the trunk fail, one of the standby ports will automatically be activated to replace it.

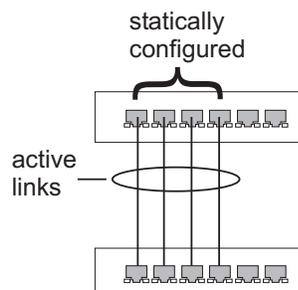
Command Usage

Besides balancing the load across each port in the trunk, the other ports provide redundancy by taking over the load if a port in the trunk fails. However, before making any physical connections between devices, use the web interface or CLI to specify the trunk on the devices at both ends. When using a trunk, take note of the following points:

- ◆ Finish configuring port trunks before you connect the corresponding network cables between switches to avoid creating a loop.
- ◆ You can create up to 27/16 trunks on a switch, with up to 54/32 ports per trunk.
- ◆ The ports at both ends of a connection must be configured as trunk ports.
- ◆ When configuring static trunks on switches of different types, they must be compatible with the Cisco EtherChannel standard.
- ◆ The ports at both ends of a trunk must be configured in an identical manner, including communication mode (i.e., speed, duplex mode and flow control), VLAN assignments, and CoS settings.
- ◆ Trunk groups are limited to either all 10G ports or all 40G ports. When using an LAG composed of all 10G ports, different transceiver types may be used as long as the speed of each member port is the same.
- ◆ All the ports in a trunk have to be treated as a whole when moved from/to, added or deleted from a VLAN.
- ◆ STP, VLAN, and IGMP settings can only be made for the entire trunk.

Configuring a Static Trunk Use the Interface > Trunk > Static page to create a trunk, assign member ports, and configure the connection parameters.

Figure 48: Configuring Static Trunks



Command Usage

- ◆ When configuring static trunks, you may not be able to link switches of different types, depending on the vendor's implementation. However, note that the static trunks on this switch are Cisco EtherChannel compatible.
- ◆ To avoid creating a loop in the network, be sure you add a static trunk via the configuration interface before connecting the ports, and also disconnect the ports before removing a static trunk via the configuration interface.

Parameters

These parameters are displayed:

- ◆ **Trunk ID** – Trunk identifier. (1-27)
- ◆ **Member** – The initial trunk member. Use the Add Member page to configure additional members.
 - **Unit** – Unit identifier. (Range: 1)
 - **Port** – Port identifier. (Range: 1- 32/54)

Web Interface

To create a static trunk:

1. Click Interface, Trunk, Static.
2. Select Configure Trunk from the Step list.
3. Select Add from the Action list.
4. Enter a trunk identifier.
5. Set the unit and port for the initial trunk member.
6. Click Apply.

Figure 49: Creating Static Trunks

Interface > Trunk > Static

Step: 1. Configure Trunk ▼ Action: Add ▼

Trunk ID (1-27) 1

Member Unit 1 ▼ Port 1 ▼

Apply Revert

To add member ports to a static trunk:

1. Click Interface, Trunk, Static.
2. Select Configure Trunk from the Step list.
3. Select Add Member from the Action list.
4. Select a trunk identifier.
5. Set the unit and port for an additional trunk member.
6. Click Apply.

Figure 50: Adding Static Trunks Members

To configure connection parameters for a static trunk:

1. Click Interface, Trunk, Static.
2. Select Configure General from the Step list.
3. Select Configure from the Action list.
4. Modify the required interface settings. (Refer to [“Configuring by Port List” on page 104](#) for a description of the parameters.)
5. Click Apply.

Figure 51: Configuring Connection Parameters for a Static Trunk

Trunk	Type	Name	Admin	Autonegotiation	Speed Duplex	Flow Control	MTU Size (1500-12288)	Link Up Down Trap
1	10GBASE SFP+		<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> Enabled <input checked="" type="checkbox"/> 10Gf <input type="checkbox"/> Sym	10Gfull	<input type="checkbox"/> Enabled	1518	<input checked="" type="checkbox"/> Enabled

To display trunk connection parameters:

1. Click Interface, Trunk, Static.
2. Select Configure General from the Step list.
3. Select Show Information from the Action list.

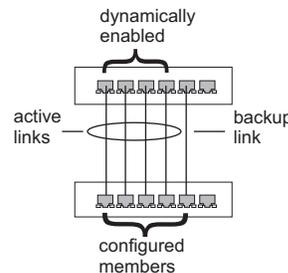
Figure 52: Showing Information for Static Trunks

Trunk	Type	Name	Admin	Oper Status	Autonegotiation	Oper Speed Duplex	Oper Flow Control	MTU Size	Link Up Down Trap
1	10GBASE SFP+		Enabled	Down	Disabled	10Gfull	None	1518	Enabled

Configuring a Dynamic Trunk

Use the Interface > Trunk > Dynamic pages to set the administrative key for an aggregation group, enable LACP on a port, configure protocol parameters for local and partner ports, or to set Ethernet connection parameters.

Figure 53: Configuring Dynamic Trunks



Command Usage

- ◆ To avoid creating a loop in the network, be sure you enable LACP before connecting the ports, and also disconnect the ports before disabling LACP.
- ◆ If the target switch has also enabled LACP on the connected ports, the trunk will be activated automatically.
- ◆ A trunk formed with another switch using LACP will automatically be assigned the next available trunk ID.
- ◆ If more than eight ports attached to the same target switch have LACP enabled, the additional ports will be placed in standby mode, and will only be enabled if one of the active links fails.
- ◆ All ports on both ends of an LACP trunk must be configured for full duplex, and auto-negotiation.

- ◆ Ports are only allowed to join the same Link Aggregation Group (LAG) if (1) the LACP port system priority matches, (2) the LACP port admin key matches, and (3) the LAG admin key matches (if configured). However, if the LAG admin key is set, then the port admin key must be set to the same value for a port to be allowed to join that group.



Note: If the LACP admin key is not set when a channel group is formed (i.e., it has a null value of 0), the operational value of this key is set to the same value as the port admin key used by the interfaces that joined the group (see the “show lacp internal” command in the *CLI Reference Guide*).

Parameters

These parameters are displayed:

Configure Aggregator

- ◆ **Admin Key** – LACP administration key is used to identify a specific link aggregation group (LAG) during local LACP setup on the switch. (Range: 0-65535)

If the port channel admin key is not set when a channel group is formed (i.e., it has the null value of 0), this key is set to the same value as the port admin key (see *Configure Aggregation Port - Actor/Partner*) used by the interfaces that joined the group. Note that when the LAG is no longer used, the port channel admin key is reset to 0.

If the port channel admin key is set to a non-default value, the operational key is based upon LACP PDUs received from the partner, and the channel admin key is reset to the default value. The trunk identifier will also be changed by this process.

- ◆ **Timeout Mode** – The timeout to wait for the next LACP data unit (LACPDU):
 - **Long Timeout** – Specifies a slow timeout of 90 seconds. (This is the default setting.)
 - **Short Timeout** – Specifies a fast timeout of 3 seconds.

The timeout is set in the LACP timeout bit of the Actor State field in transmitted LACPDU. When the partner switch receives an LACPDU set with a short timeout from the actor switch, the partner adjusts the transmit LACPDU interval to 1 second. When it receives an LACPDU set with a long timeout from the actor, it adjusts the transmit LACPDU interval to 30 seconds.

If the actor does not receive an LACPDU from its partner before the configured timeout expires, the partner port information will be deleted from the LACP group.

When a dynamic port-channel member leaves a port-channel, the default timeout value will be restored on that port.

When a dynamic port-channel is torn down, the configured timeout value will be retained. When the dynamic port-channel is constructed again, that timeout value will be used.

Configure Aggregation Port - General

- ◆ **Port** – Port identifier. (Range: 1-32/54)
- ◆ **LACP Status** – Enables or disables LACP on a port.

Configure Aggregation Port - Actor/Partner

- ◆ **Port** – Port number. (Range: 1-32/54)
- ◆ **Admin Key** – The LACP administration key must be set to the same value for ports that belong to the same LAG. (Range: 0-65535; Default – Actor: 1, Partner: 0)

~~By default, the Actor Admin Key is determined by port's link speed, and copied to Oper Key. The Partner Admin Key is assigned to zero, and the Oper Key is set based upon LACP PDUs received from the Partner.~~

Once the remote side of a link has been established, LACP operational settings are already in use on that side. Configuring LACP settings for the partner only applies to its administrative state, not its operational state.



Note: Configuring the partner admin-key does not affect remote or local switch operation. The local switch just records the partner admin-key for user reference.

By default, the actor's operational key is determined by port's link speed (1000f - 4, 100f - 3, 10f - 2), and copied to the admin key.

- ◆ **System Priority** – LACP system priority is used to determine link aggregation group (LAG) membership, and to identify this device to other switches during LAG negotiations. (Range: 0-65535; Default: 32768)
System priority is combined with the switch's MAC address to form the LAG identifier. This identifier is used to indicate a specific LAG during LACP negotiations with other systems.
- ◆ **Port Priority** – If a link goes down, LACP port priority is used to select a backup link. (Range: 0-65535; Default: 32768)
 - Setting a lower value indicates a higher effective priority.
 - If an active port link goes down, the backup port with the highest priority is selected to replace the downed link. However, if two or more ports have the same LACP port priority, the port with the lowest physical port number will be selected as the backup port.
 - If an LAG already exists with the maximum number of allowed port members, and LACP is subsequently enabled on another port using a

higher priority than an existing member, the newly configured port will replace an existing port member that has a lower priority.



Note: Configuring LACP settings for a port only applies to its administrative state, not its operational state, and will only take effect the next time an aggregate link is established with that port.

Note: Configuring the port partner sets the remote side of an aggregate link; i.e., the ports on the attached device. The command attributes have the same meaning as those used for the port actor.

Web Interface

To configure the admin key for a dynamic trunk:

1. Click Interface, Trunk, Dynamic.
2. Select Configure Aggregator from the Step list.
3. Set the Admin Key and timeout mode for the required LACP group.
4. Click Apply.

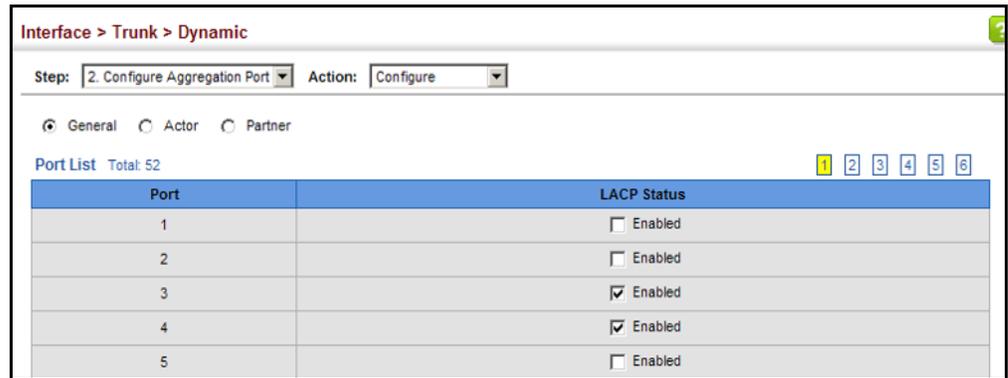
Figure 54: Configuring the LACP Aggregator Admin Key

Trunk	Admin Key (0-65535)	Timeout Mode
1	0	Long Timeout
2	0	Long Timeout
3	0	Long Timeout
4	0	Long Timeout
5	0	Long Timeout

To enable LACP for a port:

1. Click Interface, Trunk, Dynamic.
2. Select Configure Aggregation Port from the Step list.
3. Select Configure from the Action list.
4. Click General.
5. Enable LACP on the required ports.
6. Click Apply.

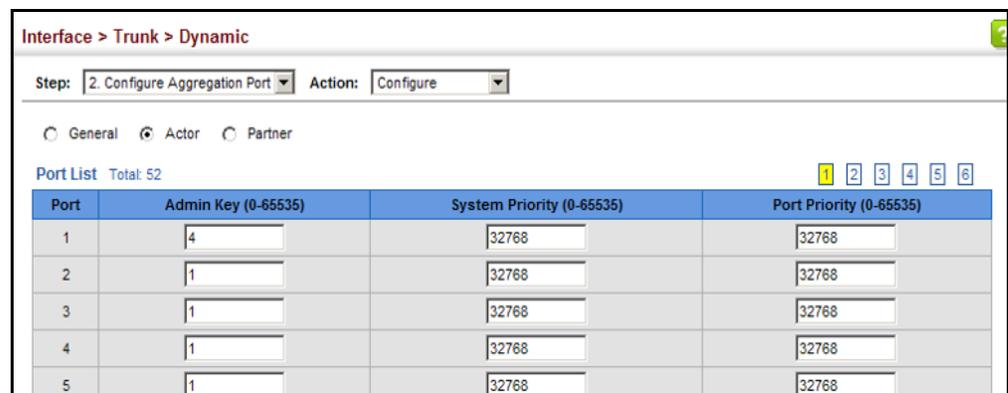
Figure 55: Enabling LACP on a Port



To configure LACP parameters for group members:

1. Click Interface, Trunk, Dynamic.
2. Select Configure Aggregation Port from the Step list.
3. Select Configure from the Action list.
4. Click Actor or Partner.
5. Configure the required settings.
6. Click Apply.

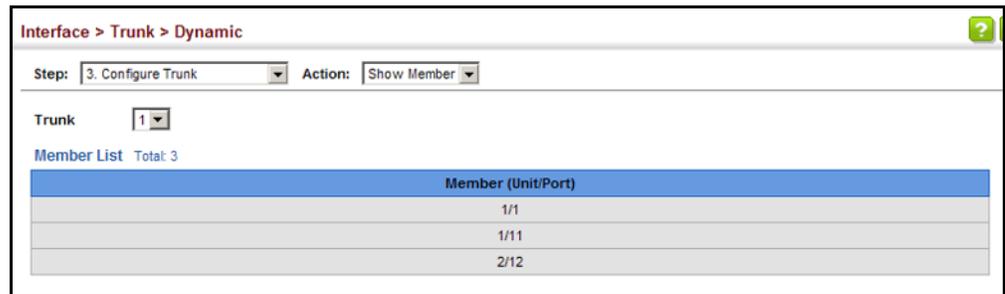
Figure 56: Configuring LACP Parameters on a Port



To show the active members of a dynamic trunk:

1. Click Interface, Trunk, Dynamic.
2. Select Configure Trunk from the Step List.
3. Select Show Member from the Action List.
4. Select a Trunk.

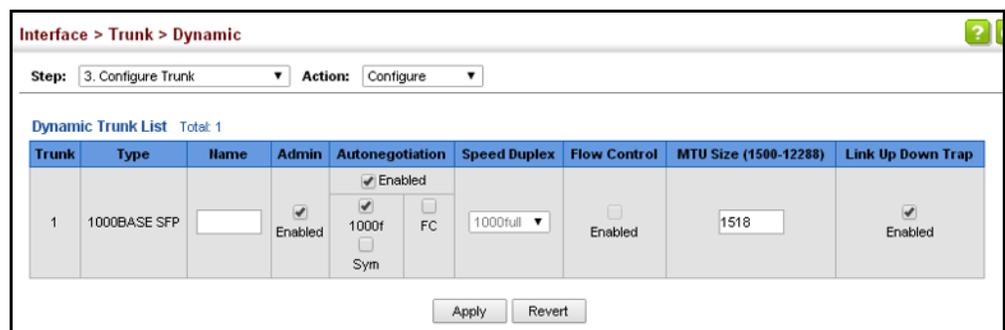
Figure 57: Showing Members of a Dynamic Trunk



To configure connection parameters for a dynamic trunk:

1. Click Interface, Trunk, Dynamic.
2. Select Configure Trunk from the Step List.
3. Select Configure from the Action List.
4. Modify the required interface settings. (See [“Configuring by Port List” on page 104](#) for a description of the interface settings.)
5. Click Apply.

Figure 58: Configuring Connection Settings for Dynamic Trunks



To display connection parameters for a dynamic trunk:

1. Click Interface, Trunk, Dynamic.
2. Select Configure Trunk from the Step List.
3. Select Show from the Action List.

Figure 59: Displaying Connection Parameters for Dynamic Trunks

Interface > Trunk > Dynamic									
Step: 3. Configure Trunk		Action: Show							
Dynamic Trunk List Total: 1									
Trunk	Type	Name	Admin	Oper Status	Autonegotiation	Oper Speed Duplex	Oper Flow Control	MTU Size	Link Up Down Trap
1	1000BASE SFP		Enabled	Up	Enabled	1000full	None	1518	Enabled

Displaying LACP Port Counters Use the Interface > Trunk > Dynamic (Configure Aggregation Port - Show Information - Counters) page to display statistics for LACP protocol messages.

Parameters

These parameters are displayed:

Table 7: LACP Port Counters

Parameter	Description
LACPDUs Sent	Number of valid LACPDUs transmitted from this channel group.
LACPDUs Received	Number of valid LACPDUs received on this channel group.
Marker Sent	Number of valid Marker PDUs transmitted from this channel group.
Marker Received	Number of valid Marker PDUs received by this channel group.
Marker Unknown Pkts	Number of frames received that either (1) Carry the Slow Protocols Ethernet Type value, but contain an unknown PDU, or (2) are addressed to the Slow Protocols group MAC Address, but do not carry the Slow Protocols Ethernet Type.
Marker Illegal Pkts	Number of frames that carry the Slow Protocols Ethernet Type value, but contain a badly formed PDU or an illegal value of Protocol Subtype.

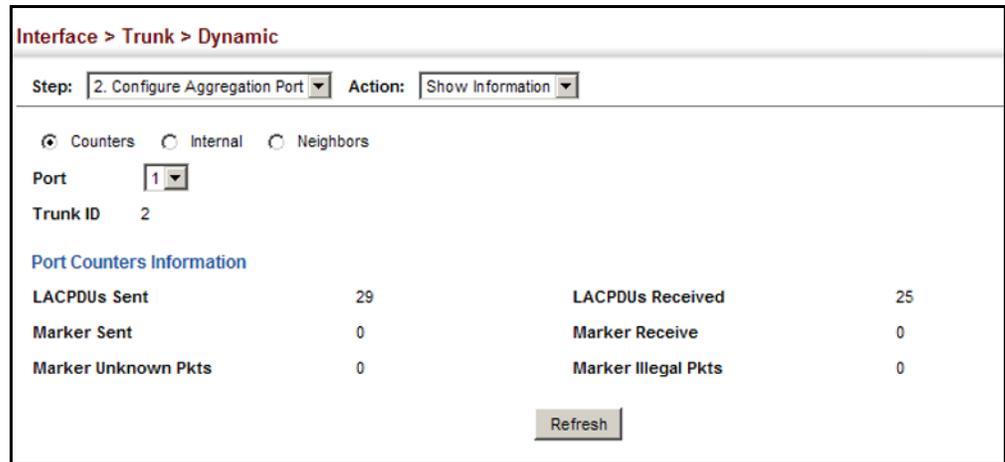
Web Interface

To display LACP port counters:

1. Click Interface, Trunk, Dynamic.
2. Select Configure Aggregation Port from the Step list.
3. Select Show Information from the Action list.
4. Click Counters.

5. Select a group member from the Port list.

Figure 60: Displaying LACP Port Counters



Displaying LACP Settings and Status for the Local Side

Use the Interface > Trunk > Dynamic (Configure Aggregation Port - Show Information - Internal) page to display the configuration settings and operational state for the local side of a link aggregation.

Parameters

These parameters are displayed:

Table 8: LACP Internal Configuration Information

Parameter	Description
LACP System Priority	LACP system priority assigned to this port channel.
LACP Port Priority	LACP port priority assigned to this interface within the channel group.
Admin Key	Current administrative value of the key for the aggregation port.
Oper Key	Current operational value of the key for the aggregation port.
LACPDU Interval	Number of seconds before invalidating received LACPDU information.
Admin State, Oper State	Administrative or operational values of the actor's state parameters: <ul style="list-style-type: none"> Expired – The actor's receive machine is in the expired state. Defaulted – The actor's receive machine is using defaulted operational partner information, administratively configured for the partner. Distributing – If false, distribution of outgoing frames on this link is disabled; i.e., distribution is currently disabled and is not expected to be enabled in the absence of administrative changes or changes in received protocol information. Collecting – Collection of incoming frames on this link is enabled; i.e., collection is currently enabled and is not expected to be disabled in the absence of administrative changes or changes in received protocol information. Synchronization – The System considers this link to be IN_SYNC; i.e., it has been allocated to the correct Link Aggregation Group, the group has been associated with a compatible Aggregator, and the identity of the Link Aggregation Group is consistent with the System ID and operational Key information transmitted.

Table 8: LACP Internal Configuration Information (Continued)

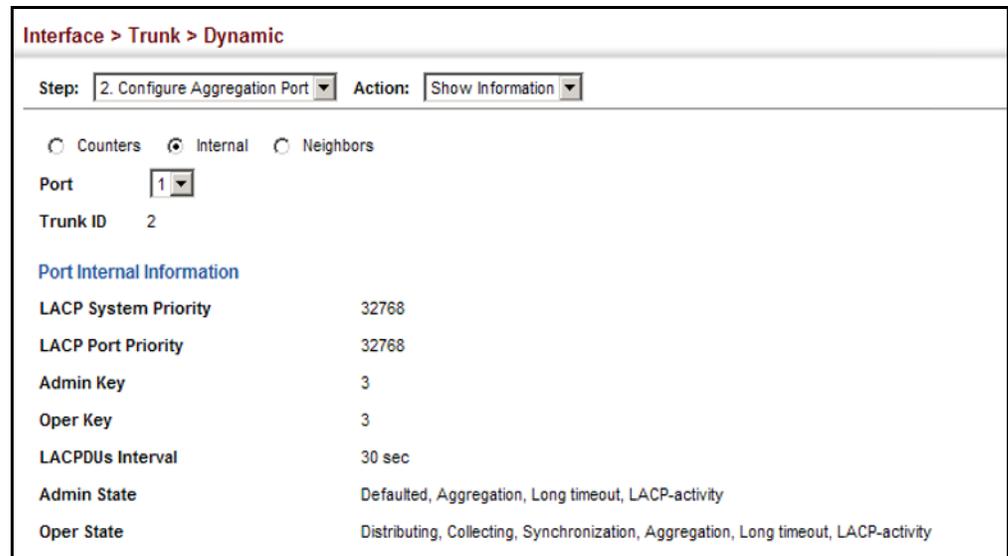
Parameter	Description
	<ul style="list-style-type: none"> ◆ Aggregation – The system considers this link to be aggregatable; i.e., a potential candidate for aggregation. ◆ Long timeout – Periodic transmission of LACPDUs uses a slow transmission rate. ◆ LACP-Activity – Activity control value with regard to this link. (0: Passive; 1: Active)

Web Interface

To display LACP settings and status for the local side:

1. Click Interface, Trunk, Dynamic.
2. Select Configure Aggregation Port from the Step list.
3. Select Show Information from the Action list.
4. Click Internal.
5. Select a group member from the Port list.

Figure 61: Displaying LACP Port Internal Information



Displaying LACP Settings and Status for the Remote Side

Use the Interface > Trunk > Dynamic (Configure Aggregation Port - Show Information - Neighbors) page to display the configuration settings and operational state for the remote side of a link aggregation.

Parameters

These parameters are displayed:

Table 9: LACP Remote Device Configuration Information

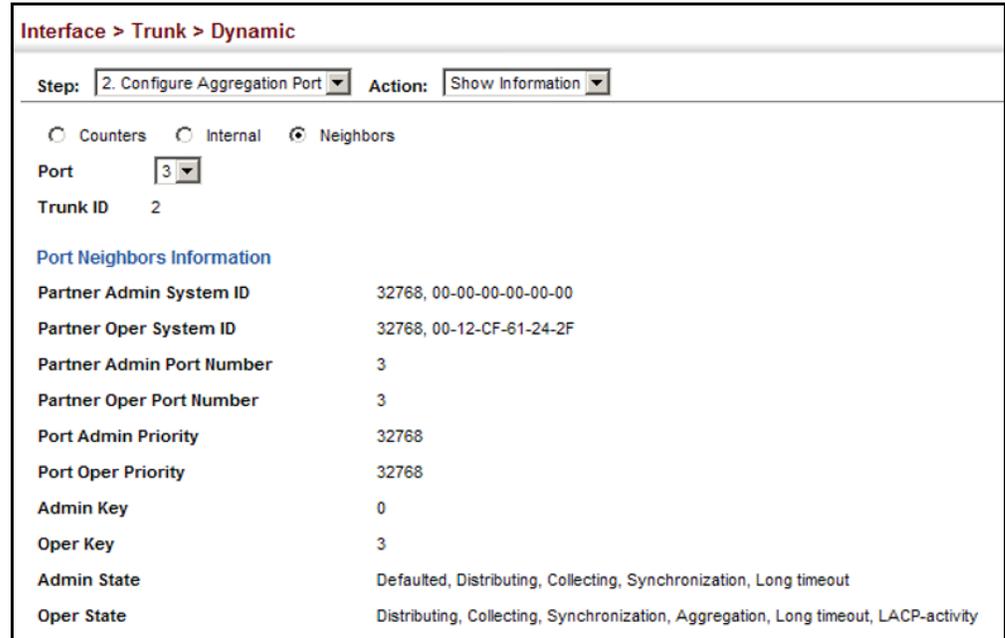
Parameter	Description
Partner Admin System ID	LAG partner's system ID assigned by the user.
Partner Oper System ID	LAG partner's system ID assigned by the LACP protocol.
Partner Admin Port Number	Current administrative value of the port number for the protocol Partner.
Partner Oper Port Number	Operational port number assigned to this aggregation port by the port's protocol partner.
Port Admin Priority	Current administrative value of the port priority for the protocol partner.
Port Oper Priority	Priority value assigned to this aggregation port by the partner.
Admin Key	Current administrative value of the Key for the protocol partner.
Oper Key	Current operational value of the Key for the protocol partner.
Admin State	Administrative values of the partner's state parameters. (See preceding table.)
Oper State	Operational values of the partner's state parameters. (See preceding table.)

Web Interface

To display LACP settings and status for the remote side:

1. Click Interface, Trunk, Dynamic.
2. Select Configure Aggregation Port from the Step list.
3. Select Show Information from the Action list.
4. Click Internal.
5. Select a group member from the Port list.

Figure 62: Displaying LACP Port Remote Information



Configuring Load Balancing Use the Interface > Trunk > Load Balance page to set the load-distribution method used among ports in aggregated links.

Command Usage

- ◆ This command applies to all static and dynamic trunks on the switch.
- ◆ To ensure that the switch traffic load is distributed evenly across all links in a trunk, select the source and destination addresses used in the load-balance calculation to provide the best result for trunk connections:
 - **Destination IP Address:** All traffic with the same destination IP address is output on the same link in a trunk. This mode works best for switch-to-router trunk links where traffic through the switch is destined for many different hosts. Do not use this mode for switch-to-server trunk links where the destination IP address is the same for all traffic.
 - **Destination MAC Address:** All traffic with the same destination MAC address is output on the same link in a trunk. This mode works best for switch-to-switch trunk links where traffic through the switch is destined for many different hosts. Do not use this mode for switch-to-router trunk links where the destination MAC address is the same for all traffic.
 - **Source and Destination IP Address:** All traffic with the same source and destination IP address is output on the same link in a trunk. This mode works best for switch-to-router trunk links where traffic through the switch is received from and destined for many different hosts.

- **Source and Destination MAC Address:** All traffic with the same source and destination MAC address is output on the same link in a trunk. This mode works best for switch-to-switch trunk links where traffic through the switch is received from and destined for many different hosts.
- **Source IP Address:** All traffic with the same source IP address is output on the same link in a trunk. This mode works best for switch-to-router or switch-to-server trunk links where traffic through the switch is received from many different hosts.
- **Source MAC Address:** All traffic with the same source MAC address is output on the same link in a trunk. This mode works best for switch-to-switch trunk links where traffic through the switch is received from many different hosts.

Parameters

These parameters are displayed for the load balance mode:

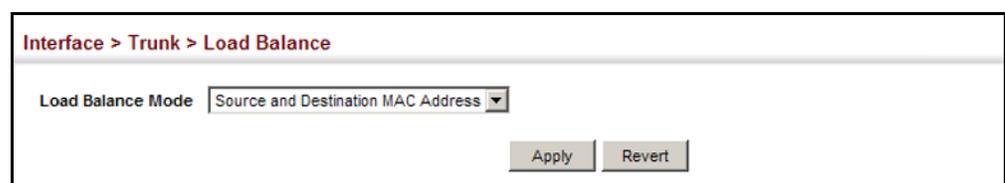
- ◆ **Destination IP Address** - Load balancing based on destination IP address.
- ◆ **Destination MAC Address** - Load balancing based on destination MAC address.
- ◆ **Source and Destination IP Address** - Load balancing based on source and destination IP address. (This is the default setting.)
- ◆ **Source and Destination MAC Address** - Load balancing based on source and destination MAC address.
- ◆ **Source IP Address** - Load balancing based on source IP address.
- ◆ **Source MAC Address** - Load balancing based on source MAC address.

Web Interface

To display the load-distribution method used by ports in aggregated links:

1. Click Interface, Trunk, Load Balance.
2. Select the required method from the Load Balance Mode list.
3. Click Apply.

Figure 63: Configuring Load Balancing



Traffic Segmentation

If tighter security is required for passing traffic from different clients through downlink ports on the local network and over uplink ports to the service provider, port-based traffic segmentation can be used to isolate traffic for individual clients. Data traffic on downlink ports is only forwarded to, and from, uplink ports

Traffic belonging to each client is isolated to the allocated downlink ports. But the switch can be configured to either isolate traffic passing across a client's allocated uplink ports from the uplink ports assigned to other clients, or to forward traffic through the uplink ports used by other clients, allowing different clients to share access to their uplink ports where security is less likely to be compromised.

Enabling Traffic Segmentation

Use the Interface > Traffic Segmentation (Configure Global) page to enable traffic segmentation.

Parameters

These parameters are displayed:

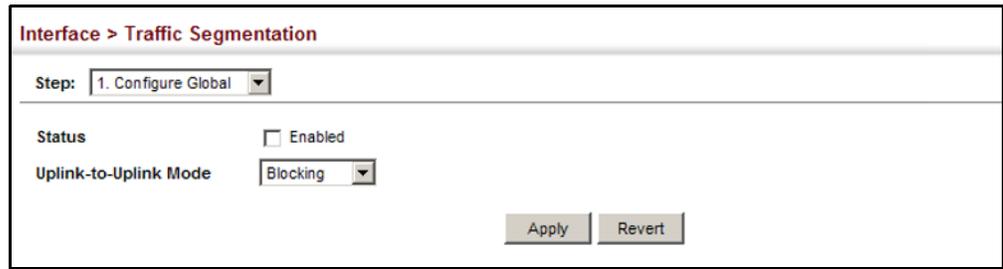
- ◆ **Status** – Enables port-based traffic segmentation. (Default: Disabled)
- ◆ **Uplink-to-Uplink Mode** – Specifies whether or not traffic can be forwarded between uplink ports assigned to different client sessions.
 - **Blocking** – Blocks traffic between uplink ports assigned to different sessions.
 - **Forwarding** – Forwards traffic between uplink ports assigned to different sessions.

Web Interface

To enable traffic segmentation:

1. Click Interface, Traffic Segmentation.
2. Select Configure Global from the Step list.
3. Mark the Status check box, and set the required uplink-to-uplink mode.
4. Click Apply.

Figure 64: Enabling Traffic Segmentation



Configuring Uplink and Downlink Ports

Use the Interface > Traffic Segmentation (Configure Session) page to assign the downlink and uplink ports to use in the segmented group. Ports designated as downlink ports can not communicate with any other ports on the switch except for the uplink ports. Uplink ports can communicate with any other ports on the switch and with any designated downlink ports.

Command Usage

- ◆ When traffic segmentation is enabled, the forwarding state for the uplink and downlink ports assigned to different client sessions is shown below.

Table 10: Traffic Segmentation Forwarding

Destination Source	Session #1 Downlinks	Session #1 Uplinks	Session #2 Downlinks	Session #2 Uplinks	Normal Ports
Session #1 Downlink Ports	Blocking	Forwarding	Blocking	Blocking	Blocking
Session #1 Uplink Ports	Forwarding	Forwarding	Blocking	Blocking/ Forwarding*	Forwarding
Session #2 Downlink Ports	Blocking	Blocking	Blocking	Forwarding	Blocking
Session #2 Uplink Ports	Blocking	Blocking/ Forwarding*	Forwarding	Forwarding	Forwarding
Normal Ports	Forwarding	Forwarding	Forwarding	Forwarding	Forwarding

* The forwarding state for uplink-to-uplink ports is configured on the Configure Global page (see [page 142](#)).

- ◆ When traffic segmentation is disabled, all ports operate in normal forwarding mode based on the settings specified by other functions such as VLANs and spanning tree protocol.
- ◆ A port cannot be configured in both an uplink and downlink list.
- ◆ A port can only be assigned to one traffic-segmentation session.
- ◆ A downlink port can only communicate with an uplink port in the same session. Therefore, if an uplink port is not configured for a session, the assigned downlink ports will not be able to communicate with any other ports.

- ◆ If a downlink port is not configured for the session, the assigned uplink ports will operate as normal ports.

Parameters

These parameters are displayed:

- ◆ **Session ID** – Traffic segmentation session. (Range: 1-4)
- ◆ **Direction** – Adds an interface to the segmented group by setting the direction to uplink or downlink. (Default: Uplink)
- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Port** – Port Identifier. (Range: 1-32/54)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-27)

Web Interface

To configure the members of the traffic segmentation group:

1. Click Interface, Traffic Segmentation.
2. Select Configure Session from the Step list.
3. Select Add from the Action list.
4. Enter the session ID, set the direction to uplink or downlink, and select the interface to add.
5. Click Apply.

Figure 65: Configuring Members for Traffic Segmentation

The screenshot shows a web interface titled "Interface > Traffic Segmentation". At the top, there are two dropdown menus: "Step: 2. Configure Session" and "Action: Add". Below these, there are three main sections: "Session ID (1-4)" with a text input field; "Direction" with a dropdown menu set to "Uplink"; and "Interface" with two radio button options: "Port (1-54)" (selected) and "Trunk (1-27)". Each radio button option has two adjacent text input fields. At the bottom right, there are two buttons: "Apply" and "Revert".

To show the members of the traffic segmentation group:

1. Click Interface, Traffic Segmentation.
2. Select Configure Session from the Step list.

3. Select Show from the Action list.

Figure 66: Showing Traffic Segmentation Members

Interface > Traffic Segmentation

Step: 2. Configure Session Action: Show

Session List Total: 2

<input type="checkbox"/>	Session ID	Direction	Interface
<input type="checkbox"/>	1	Uplink	Unit 1 / Port 1
<input type="checkbox"/>	1	Downlink	Unit 1 / Port 2

Delete Revert

5

VLAN Configuration

This chapter includes the following topics:

- ◆ [IEEE 802.1Q VLANs](#) – Configures static and dynamic VLANs.
- ◆ [IEEE 802.1Q Tunneling](#) – Configures QinQ tunneling to maintain customer-specific VLAN and Layer 2 protocol configurations across a service provider network, even when different customers use the same internal VLAN IDs.

IEEE 802.1Q VLANs

In large networks, routers are used to isolate broadcast traffic for each subnet into separate domains. This switch provides a similar service at Layer 2 by using VLANs to organize any group of network nodes into separate broadcast domains. VLANs confine broadcast traffic to the originating group, and can eliminate broadcast storms in large networks. This also provides a more secure and cleaner network environment.

An IEEE 802.1Q VLAN is a group of ports that can be located anywhere in the network, but communicate as though they belong to the same physical segment.

VLANs help to simplify network management by allowing you to move devices to a new VLAN without having to change any physical connections. VLANs can be easily organized to reflect departmental groups (such as Marketing or R&D), usage groups (such as e-mail), or multicast groups (used for multimedia applications such as video conferencing).

VLANs provide greater network efficiency by reducing broadcast traffic, and allow you to make network changes without having to update IP addresses or IP subnets. VLANs inherently provide a high level of network security since traffic must pass through a configured Layer 3 link to reach a different VLAN.

This switch supports the following VLAN features:

- ◆ Up to 4094 VLANs based on the IEEE 802.1Q standard
- ◆ Distributed VLAN learning across multiple switches using explicit or implicit tagging
- ◆ Port overlapping, allowing a port to participate in multiple VLANs
- ◆ End stations can belong to multiple VLANs

- ◆ Passing traffic between VLAN-aware and VLAN-unaware devices
- ◆ Priority tagging

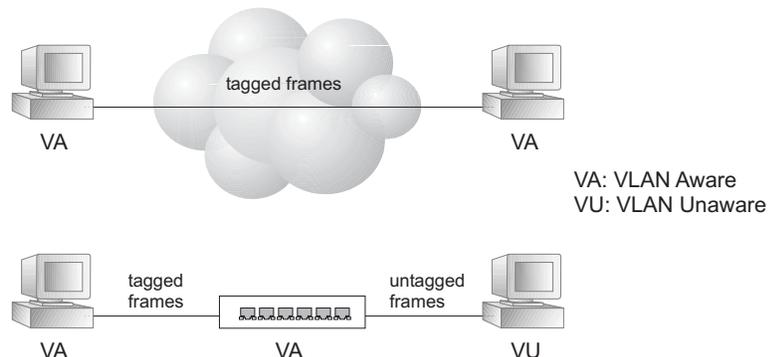
Assigning Ports to VLANs

Before enabling VLANs for the switch, you must first assign each port to the VLAN group(s) in which it will participate. By default all ports are assigned to VLAN 1 as untagged ports. Add a port as a tagged port if you want it to carry traffic for one or more VLANs, and any intermediate network devices or the host at the other end of the connection supports VLANs. Then manually assign ports on the other VLAN-aware network devices along the path that will carry this traffic to the same VLAN(s). However, if you want a port on this switch to participate in one or more VLANs, but none of the intermediate network devices nor the host at the other end of the connection supports VLANs, then you should add this port to the VLAN as an untagged port.



Note: VLAN-tagged frames can pass through VLAN-aware or VLAN-unaware network interconnection devices, but the VLAN tags should be stripped off before passing it on to any end-node host that does not support VLAN tagging.

Figure 67: VLAN Compliant and VLAN Non-compliant Devices



VLAN Classification – When the switch receives a frame, it classifies the frame in one of two ways. If the frame is untagged, the switch assigns the frame to an associated VLAN (based on the default VLAN ID of the receiving port). But if the frame is tagged, the switch uses the tagged VLAN ID to identify the port broadcast domain of the frame.

Port Overlapping – Port overlapping can be used to allow access to commonly shared network resources among different VLAN groups, such as file servers or printers. Note that if you implement VLANs which do not overlap, but still need to communicate, you can connect them by enabled routing on this switch.

Untagged VLANs – Untagged VLANs are typically used to reduce broadcast traffic and to increase security. A group of network users assigned to a VLAN form a broadcast domain that is separate from other VLANs configured on the switch. Packets are forwarded only between ports that are designated for the same VLAN.

Untagged VLANs can be used to manually isolate user groups or subnets. However, you should use IEEE 802.3 tagged VLANs whenever possible to automate VLAN registration.

Forwarding Tagged/Untagged Frames

If you want to create a small port-based VLAN for devices attached directly to a single switch, you can assign ports to the same untagged VLAN. However, to participate in a VLAN group that crosses several switches, you should create a VLAN for that group and enable tagging on all ports.

Ports can be assigned to multiple tagged or untagged VLANs. Each port on the switch is therefore capable of passing tagged or untagged frames. When forwarding a frame from this switch along a path that contains any VLAN-aware devices, the switch should include VLAN tags. When forwarding a frame from this switch along a path that does not contain any VLAN-aware devices (including the destination host), the switch must first strip off the VLAN tag before forwarding the frame. When the switch receives a tagged frame, it will pass this frame onto the VLAN(s) indicated by the frame tag. However, when this switch receives an untagged frame from a VLAN-unaware device, it first decides where to forward the frame, and then inserts a VLAN tag reflecting the ingress port's default VID.

Configuring VLAN Groups

Use the VLAN > Static (Add) page to create or remove VLAN groups, or set administrative status. To propagate information about VLAN groups used on this switch to external network devices, you must specify a VLAN ID for each of these groups.

Parameters

These parameters are displayed:

Add

- ◆ **VLAN ID** – ID of VLAN or range of VLANs (1-4094).
Up to 4094 VLAN groups can be defined. VLAN 1 is the default untagged VLAN.
- ◆ **Status** – Enables or disables the specified VLAN.
- ◆ **Remote VLAN** – Reserves this VLAN for RSPAN (see [“Configuring Remote Port Mirroring”](#) on page 110).
- ◆ **L3 Interface** – Sets the interface to support Layer 3 configuration, and reserves memory space required to maintain additional information about this interface type. This parameter must be enabled before you can assign an IP address to a VLAN (see [“Setting the Switch's IP Address \(IP Version 4\)”](#) on page 481).

Modify

- ◆ **VLAN ID** – ID of configured VLAN (1-4094).

- ◆ **VLAN Name** – Name of the VLAN (1 to 32 characters).
- ◆ **Status** – Enables or disables the specified VLAN.
- ◆ **L3 Interface** – Sets the interface to support Layer 3 configuration, and reserves memory space required to maintain additional information about this interface type. This parameter must be enabled before you can assign an IP address to a VLAN.

Show

- ◆ **VLAN ID** – ID of configured VLAN.
- ◆ **VLAN Name** – Name of the VLAN.
- ◆ **Status** – Operational status of configured VLAN.
- ◆ **Remote VLAN** – Reserves this VLAN for RSPAN (see [“Configuring Remote Port Mirroring” on page 110](#)).
- ◆ **L3 Interface** – Shows if the interface supports Layer 3 configuration.

Web Interface

To create VLAN groups:

1. Click VLAN, Static.
2. Select Add from the Action list.
3. Enter a VLAN ID or range of IDs.
4. Click Status to configure the VLAN as operational.
5. Enable the L3 Interface field to specify that a VLAN will be used as a Layer 3 interface.
6. Specify whether the VLANs are to be used for remote port mirroring.
7. Click Apply.

Figure 68: Creating Static VLANs

The screenshot shows the 'VLAN > Static' configuration page. At the top, the 'Action' dropdown is set to 'Add'. Below this, there are three main configuration sections: 'VLAN ID (1-4094)' with a text input field containing '2' and a hyphen followed by an empty field; 'Status' with a checked checkbox labeled 'Enabled'; and 'Remote VLAN' with an unchecked checkbox labeled 'Enabled'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

To modify the configuration settings for VLAN groups:

1. Click VLAN, Static.
2. Select Modify from the Action list.
3. Select the identifier of a configured VLAN.
4. Modify the VLAN name, operational status, or Layer 3 Interface status as required.
5. Enable the L3 Interface field to specify that a VLAN will be used as a Layer 3 interface.
6. Click Apply.

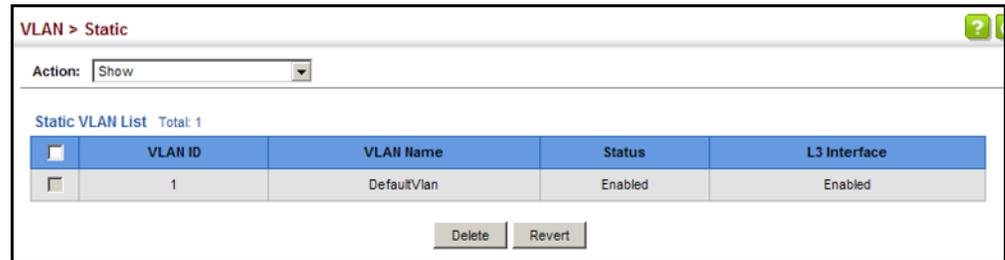
Figure 69: Modifying Settings for Static VLANs

The screenshot shows the 'VLAN > Static' configuration page with the 'Action' dropdown set to 'Modify'. The 'VLAN ID (1-4094)' is now a dropdown menu showing '1'. The 'VLAN Name' is a text input field containing 'DefaultVlan'. The 'Status' checkbox is checked and labeled 'Enabled'. The 'L3 Interface' checkbox is also checked and labeled 'Enabled'. The 'Apply' and 'Revert' buttons are at the bottom right.

To show the configuration settings for VLAN groups:

1. Click VLAN, Static.
2. Select Show from the Action list.

Figure 70: Showing Static VLANs



Adding Static Members to VLANs

Use the VLAN > Static (Edit Member by VLAN, Edit Member by Interface, or Edit Member by Interface Range) pages to configure port members for the selected VLAN index, interface, or a range of interfaces. Use the menus for editing port members to configure the VLAN behavior for specific interfaces, including the mode of operation (Hybrid or 1Q Trunk), the default VLAN identifier (PVID), accepted frame types, and ingress filtering. Assign ports as tagged if they are connected to 802.1Q VLAN compliant devices, or untagged they are not connected to any VLAN-aware devices. Or configure a port as forbidden to prevent the switch from automatically adding it to a VLAN via the GVRP protocol.

Parameters

These parameters are displayed:

Edit Member by VLAN

- ◆ **VLAN** – ID of configured VLAN (1-4094).
- ◆ **Interface** – Displays a list of ports or trunks.
 - **Port** – Port Identifier. (Range: 1-32/54)
 - **Trunk** – Trunk Identifier. (Range: 1-27)
- ◆ **Mode** – Indicates VLAN membership mode for an interface. (Default: Hybrid)
 - **Access** – Sets the port to operate as an untagged interface. The port transmits and receives untagged frames on a single VLAN only.
 - **Hybrid** – Specifies a hybrid VLAN interface. The port may transmit tagged or untagged frames.
 - **1Q Trunk** – Specifies a port as an end-point for a VLAN trunk. A trunk is a direct link between two switches, so the port transmits tagged frames that

identify the source VLAN. Note that frames belonging to the port's default VLAN (i.e., associated with the PVID) are also transmitted as tagged frames.

- ◆ **PVID** – VLAN ID assigned to untagged frames received on the interface. (Default: 1)

When using Access mode, and an interface is assigned to a new VLAN, its PVID is automatically set to the identifier for that VLAN. When using Hybrid mode, the PVID for an interface can be set to any VLAN for which it is an untagged member.

- ◆ **Acceptable Frame Type** – Sets the interface to accept all frame types, including tagged or untagged frames, or only tagged frames. When set to receive all frame types, any received frames that are untagged are assigned to the default VLAN. (Options: All, Tagged; Default: All)
- ◆ **Ingress Filtering** – Determines how to process frames tagged for VLANs for which the ingress port is not a member. (Default: Disabled)
 - Ingress filtering only affects tagged frames.
 - If ingress filtering is disabled and a port receives frames tagged for VLANs for which it is not a member, these frames will be flooded to all other ports (except for those VLANs explicitly forbidden on this port).
 - If ingress filtering is enabled and a port receives frames tagged for VLANs for which it is not a member, these frames will be discarded.
 - Ingress filtering does not affect VLAN independent BPDU frames, such as GVRP or STP. However, they do affect VLAN dependent BPDU frames, such as GMRP.
- ◆ **Membership Type** – Select VLAN membership for each interface by marking the appropriate radio button for a port or trunk:
 - **Tagged:** Interface is a member of the VLAN. All packets transmitted by the port will be tagged, that is, carry a tag and therefore carry VLAN or CoS information.
 - **Untagged:** Interface is a member of the VLAN. All packets transmitted by the port will be untagged, that is, not carry a tag and therefore not carry VLAN or CoS information. Note that an interface must be assigned to at least one group as an untagged port.
 - **None:** Interface is not a member of the VLAN. Packets associated with this VLAN will not be transmitted by the interface.



Note: VLAN 1 is the default untagged VLAN containing all ports on the switch.

Edit Member by Interface

All parameters are the same as those described under the preceding section for Edit Member by VLAN.

Edit Member by Interface Range

All parameters are the same as those described under the earlier section for Edit Member by VLAN, except for the items shown below.

- ◆ **Port Range** – Displays a list of ports. (Range: 1-32/54)
- ◆ **Trunk Range** – Displays a list of ports. (Range: 1-27)



Note: The PVID, acceptable frame type, and ingress filtering parameters for each interface within the specified range must be configured on either the Edit Member by VLAN or Edit Member by Interface page.

Web Interface

To configure static members by the VLAN index:

1. Click VLAN, Static.
2. Select Edit Member by VLAN from the Action list.
3. Set the Interface type to display as Port or Trunk.
4. Modify the settings for any interface as required.
5. Click Apply.

Figure 71: Configuring Static Members by VLAN Index

Port	Mode	PVID	Acceptable Frame Type	Ingress Filtering	Membership Type			
					Tagged	Untagged	Forbidden	None
1	Hybrid	1	All	<input type="checkbox"/> Enabled	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
2	Hybrid	1	All	<input type="checkbox"/> Enabled	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
3	Hybrid	1	All	<input type="checkbox"/> Enabled	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
4	Hybrid	1	All	<input type="checkbox"/> Enabled	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
5	Hybrid	1	All	<input type="checkbox"/> Enabled	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

To configure static members by interface:

1. Click VLAN, Static.
2. Select Edit Member by Interface from the Action list.
3. Select a port or trunk to configure.
4. Modify the settings for any interface as required.
5. Click Apply.

Figure 72: Configuring Static VLAN Members by Interface

The screenshot shows the 'VLAN > Static' configuration page. The 'Action' dropdown is set to 'Edit Member by Interface'. The 'Interface' section has 'Port' selected with '1' in the dropdown, and 'Trunk' is unselected. The 'Mode' is set to 'Hybrid', 'PVID' is '1', and 'Acceptable Frame Type' is 'All'. 'Ingress Filtering' is unchecked. Below this is the 'Static VLAN Membership List' with a total of 4 members. The table has columns for 'VLAN' and 'Membership Type' (Tagged, Untagged, Forbidden, None). Each row has radio buttons for each membership type.

VLAN	Membership Type			
	Tagged	Untagged	Forbidden	None
1	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
2	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
3	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
4	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>

Buttons for 'Apply' and 'Revert' are located at the bottom right of the table.

To configure static members by interface range:

1. Click VLAN, Static.
2. Select Edit Member by Interface Range from the Action list.
3. Set the Interface type to display as Port or Trunk.
4. Enter an interface range.
5. Modify the VLAN parameters as required. Remember that the PVID, acceptable frame type, and ingress filtering parameters for each interface within the specified range must be configured on either the Edit Member by VLAN or Edit Member by Interface page.
6. Click Apply.

Figure 73: Configuring Static VLAN Members by Interface Range

The screenshot shows a web-based configuration interface for VLANs. At the top, it says "VLAN > Static". Below that, there is an "Action:" dropdown menu set to "Edit Member by Interface Range". The main configuration area includes several fields and options: "Interface" with radio buttons for "Port" (selected) and "Trunk"; "Port Range (1-54)" with two input boxes separated by a hyphen; "Mode" with a dropdown menu set to "Hybrid"; "VLAN ID (1-4094)" with two input boxes separated by a hyphen; and "Membership Type" with radio buttons for "Tagged" (selected), "Untagged", and "None". At the bottom right of the form are "Apply" and "Revert" buttons.

IEEE 802.1Q Tunneling

IEEE 802.1Q Tunneling (QinQ) is designed for service providers carrying traffic for multiple customers across their networks. QinQ tunneling is used to maintain customer-specific VLAN and Layer 2 protocol configurations even when different customers use the same internal VLAN IDs. This is accomplished by inserting Service Provider VLAN (SPVLAN) tags into the customer's frames when they enter the service provider's network, and then stripping the tags when the frames leave the network.

A service provider's customers may have specific requirements for their internal VLAN IDs and number of VLANs supported. VLAN ranges required by different customers in the same service-provider network might easily overlap, and traffic passing through the infrastructure might be mixed. Assigning a unique range of VLAN IDs to each customer would restrict customer configurations, require intensive processing of VLAN mapping tables, and could easily exceed the maximum VLAN limit of 4096.

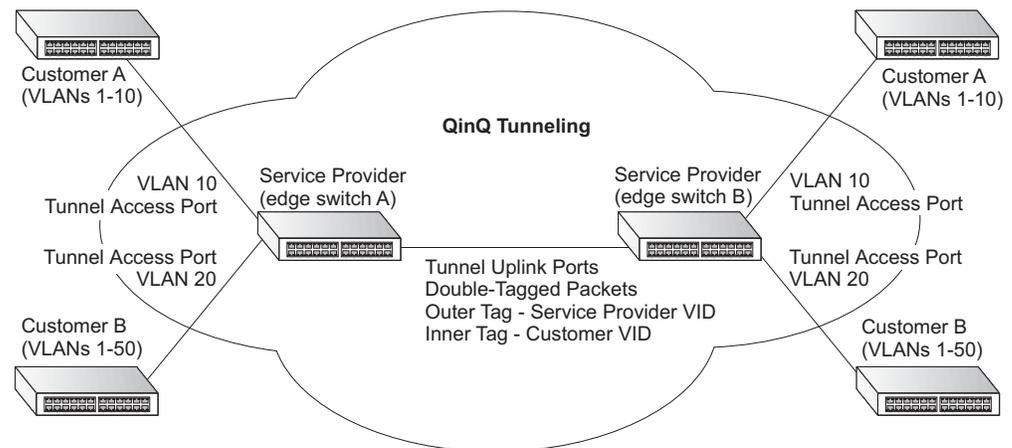
QinQ tunneling uses a single Service Provider VLAN (SPVLAN) for customers who have multiple VLANs. Customer VLAN IDs are preserved and traffic from different customers is segregated within the service provider's network even when they use the same customer-specific VLAN IDs. QinQ tunneling expands VLAN space by using a VLAN-in-VLAN hierarchy, preserving the customer's original tagged packets, and adding SPVLAN tags to each frame (also called double tagging).

A port configured to support QinQ tunneling must be set to tunnel port mode. The Service Provider VLAN (SPVLAN) ID for the specific customer must be assigned to the QinQ tunnel access port on the edge switch where the customer traffic enters the service provider's network. Each customer requires a separate SPVLAN, but this VLAN supports all of the customer's internal VLANs. The QinQ tunnel uplink port that passes traffic from the edge switch into the service provider's metro network must also be added to this SPVLAN. The uplink port can be added to multiple SPVLANs to carry inbound traffic for different customers onto the service provider's network.

When a double-tagged packet enters another trunk port in an intermediate or core switch in the service provider's network, the outer tag is stripped for packet processing. When the packet exits another trunk port on the same core switch, the same SPVLAN tag is again added to the packet.

When a packet enters the trunk port on the service provider's egress switch, the outer tag is again stripped for packet processing. However, the SPVLAN tag is not added when it is sent out the tunnel access port on the edge switch into the customer's network. The packet is sent as a normal IEEE 802.1Q-tagged frame, preserving the original VLAN numbers used in the customer's network.

Figure 74: QinQ Operational Concept



Layer 2 Flow for Packets Coming into a Tunnel Access Port

A QinQ tunnel port may receive either tagged or untagged packets. No matter how many tags the incoming packet has, it is treated as tagged packet.

The ingress process does source and destination lookups. If both lookups are successful, the ingress process writes the packet to memory. Then the egress process transmits the packet. Packets entering a QinQ tunnel port are processed in the following manner:

1. An SPVLAN tag is added to all outbound packets on the SPVLAN interface, no matter how many tags they already have. The switch constructs and inserts the outer tag (SPVLAN) into the packet based on the default VLAN ID and Tag Protocol Identifier (TPID, that is, the ether-type of the tag), unless otherwise defined as described under ["Creating CVLAN to SPVLAN Mapping Entries" on page 159](#)). The priority of the inner tag is copied to the outer tag if it is a tagged or priority tagged packet.
2. After successful source and destination lookup, the ingress process sends the packet to the switching process with two tags. If the incoming packet is untagged, the outer tag is an SPVLAN tag, and the inner tag is a dummy tag (8100 0000). If the incoming packet is tagged, the outer tag is an SPVLAN tag, and the inner tag is a CVLAN tag.

3. After packet classification through the switching process, the packet is written to memory with one tag (an outer tag) or with two tags (both an outer tag and inner tag).
4. The switch sends the packet to the proper egress port.
5. If the egress port is an untagged member of the SPVLAN, the outer tag will be stripped. If it is a tagged member, the outgoing packets will have two tags.

Layer 2 Flow for Packets Coming into a Tunnel Uplink Port

An uplink port receives one of the following packets:

- ◆ Untagged
- ◆ One tag (CVLAN or SPVLAN)
- ◆ Double tag (CVLAN + SPVLAN)

The ingress process does source and destination lookups. If both lookups are successful, the ingress process writes the packet to memory. Then the egress process transmits the packet. Packets entering a QinQ uplink port are processed in the following manner:

1. If incoming packets are untagged, the PVID VLAN native tag is added.
2. If the ether-type of an incoming packet (single or double tagged) is not equal to the TPID of the uplink port, the VLAN tag is determined to be a Customer VLAN (CVLAN) tag. The uplink port's PVID VLAN native tag is added to the packet. This outer tag is used for learning and switching packets within the service provider's network. The TPID must be configured on a per port basis, and the verification cannot be disabled.
3. If the ether-type of an incoming packet (single or double tagged) is equal to the TPID of the uplink port, no new VLAN tag is added. If the uplink port is not the member of the outer VLAN of the incoming packets, the packet will be dropped when ingress filtering is enabled. If ingress filtering is not enabled, the packet will still be forwarded. If the VLAN is not listed in the VLAN table, the packet will be dropped.
4. After successful source and destination lookups, the packet is double tagged. The switch uses the TPID of 0x8100 to indicate that an incoming packet is double-tagged. If the outer tag of an incoming double-tagged packet is equal to the port TPID and the inner tag is 0x8100, it is treated as a double-tagged packet. If a single-tagged packet has 0x8100 as its TPID, and port TPID is not 0x8100, a new VLAN tag is added and it is also treated as double-tagged packet.
5. If the destination address lookup fails, the packet is sent to all member ports of the outer tag's VLAN.

6. After packet classification, the packet is written to memory for processing as a single-tagged or double-tagged packet.
7. The switch sends the packet to the proper egress port.
8. If the egress port is an untagged member of the SPVLAN, the outer tag will be stripped. If it is a tagged member, the outgoing packet will have two tags.

Configuration Limitations for QinQ

- ◆ The native VLAN of uplink ports should not be used as the SPVLAN. If the SPVLAN is the uplink port's native VLAN, the uplink port must be an untagged member of the SPVLAN. Then the outer SPVLAN tag will be stripped when the packets are sent out. Another reason is that it causes non-customer packets to be forwarded to the SPVLAN.
- ◆ Static trunk port groups are compatible with QinQ tunnel ports as long as the QinQ configuration is consistent within a trunk port group.
- ◆ The native VLAN (VLAN 1) is not normally added to transmitted frames. Avoid using VLAN 1 as an SPVLAN tag for customer traffic to reduce the risk of misconfiguration. Instead, use VLAN 1 as a management VLAN instead of a data VLAN in the service provider network.
- ◆ There are some inherent incompatibilities between Layer 2 and Layer 3 switching:
 - Tunnel ports do not support IP Access Control Lists.
 - Layer 3 Quality of Service (QoS) and other QoS features containing Layer 3 information are not supported on tunnel ports.
 - Spanning tree bridge protocol data unit (BPDU) filtering is automatically disabled on a tunnel port.

General Configuration Guidelines for QinQ

1. Enable Tunnel Status, and set the Tag Protocol Identifier (TPID) value of the tunnel access port (in the Ethernet Type field). This step is required if the attached client is using a nonstandard 2-byte ethertype to identify 802.1Q tagged frames. The default ethertype value is 0x8100. (See [“Enabling QinQ Tunneling on the Switch” on page 158.](#))
2. Create a Service Provider VLAN, also referred to as an SPVLAN (see [“Configuring VLAN Groups” on page 147.](#))
3. Configure the QinQ tunnel access port to Access mode (see [“Adding an Interface to a QinQ Tunnel” on page 161.](#))
4. Configure the QinQ tunnel access port to join the SPVLAN as an untagged member (see [“Adding Static Members to VLANs” on page 150.](#))

5. Configure the SPVLAN ID as the native VID on the QinQ tunnel access port (see [“Adding Static Members to VLANs” on page 150](#)).
6. Configure the QinQ tunnel uplink port to Uplink mode (see [“Adding an Interface to a QinQ Tunnel” on page 161](#)).
7. Configure the QinQ tunnel uplink port to join the SPVLAN as a tagged member (see [“Adding Static Members to VLANs” on page 150](#)).

Enabling QinQ Tunneling on the Switch

Use the VLAN > Tunnel (Configure Global) page to configure the switch to operate in IEEE 802.1Q (QinQ) tunneling mode, which is used for passing Layer 2 traffic across a service provider’s metropolitan area network. You can also globally set the Tag Protocol Identifier (TPID) value of the tunnel port if the attached client is using a nonstandard 2-byte ethertype to identify 802.1Q tagged frames.

Parameters

These parameters are displayed:

- ◆ **Tunnel Status** – Sets the switch to QinQ mode. (Default: Disabled)
- ◆ **Ethernet Type** – The Tag Protocol Identifier (TPID) specifies the ethertype of incoming packets on a tunnel port. (Range: hexadecimal 0800-FFFF; Default: 8100)

Use this field to set a custom 802.1Q ethertype value for the 802.1Q Tunnel TPID. This feature allows the switch to interoperate with third-party switches that do not use the standard 0x8100 ethertype to identify 802.1Q-tagged frames. For example, if 0x1234 is set as the custom 802.1Q ethertype on a trunk port, incoming frames containing that ethertype are assigned to the VLAN contained in the tag following the ethertype field, as they would be with a standard 802.1Q trunk. Frames arriving on the port containing any other ethertype are looked upon as untagged frames, and assigned to the native VLAN of that port.

The specified ethertype only applies to ports configured in Uplink mode (see [“Adding an Interface to a QinQ Tunnel” on page 161](#)). If the port is in normal mode, the TPID is always 8100. If the port is in Access mode, received packets are processed as untagged packets.

Avoid using well-known ethertypes for the TPID unless you can eliminate all side effects. For example, setting the TPID to 0800 hexadecimal (which is used for IPv4) will interfere with management access through the web interface.

Web Interface

To enable QinQ Tunneling on the switch:

1. Click VLAN, Tunnel.
2. Select Configure Global from the Step list.

3. Enable Tunnel Status, and specify the TPID if a client attached to a tunnel port is using a non-standard ethertype to identify 802.1Q tagged frames.
4. Click Apply.

Figure 75: Enabling QinQ Tunneling

The screenshot shows a web interface for configuring a tunnel. At the top, it says 'VLAN > Tunnel'. Below that, there is a 'Step:' dropdown menu set to '1. Configure Global'. The main configuration area has two sections: 'Tunnel Status' with an unchecked checkbox labeled 'Enabled', and 'Ethernet Type' with a text input field containing '100'. Below the input field is the text '(800-FFFF, hexadecimal value)'. At the bottom right of the configuration area are two buttons: 'Apply' and 'Revert'.

Creating CVLAN to SPVLAN Mapping Entries

Use the VLAN > Tunnel (Configure Service) page to create a CVLAN to SPVLAN mapping entry.

Command Usage

- ◆ The inner VLAN tag of a customer packet entering the edge router of a service provider’s network is mapped to an outer tag indicating the service provider VLAN that will carry this traffic across the 802.1Q tunnel. By default, the outer tag is based on the default VID of the edge router’s ingress port. This process is performed in a transparent manner as described under [“IEEE 802.1Q Tunneling” on page 154](#).
- ◆ When priority bits are found in the inner tag, these are also copied to the outer tag. This allows the service provider to differentiate service based on the indicated priority and appropriate methods of queue management at intermediate nodes across the tunnel.
- ◆ Rather than relying on standard service paths and priority queuing, QinQ VLAN mapping can be used to further enhance service by defining a set of differentiated service pathways to follow across the service provider’s network for traffic arriving from specified inbound customer VLANs.
- ◆ Note that all customer interfaces should be configured as access interfaces (that is, a user-to-network interface) and service provider interfaces as uplink interfaces (that is, a network-to-network interface). Use the Configure Interface page described in the next section to set an interface to access or uplink mode.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-52)

- ◆ **Customer VLAN ID** – VLAN ID for the inner VLAN tag. (Range: 1-4094)
- ◆ **Service VLAN ID** – VLAN ID for the outer VLAN tag. (Range: 1-4094)

Web Interface

To configure a mapping entry:

1. Click VLAN, Tunnel.
2. Select Configure Service from the Step list.
3. Select Add from the Action list.
4. Select an interface from the Port list.
5. Specify the CVID to SVID mapping for packets exiting the specified port.
6. Click Apply.

Figure 76: Configuring CVLAN to SPVLAN Mapping Entries

The screenshot shows a web interface titled "VLAN > Tunnel". At the top, there are two dropdown menus: "Step:" set to "2. Configure Service" and "Action:" set to "Add". Below these, there is a "Port" dropdown menu set to "1". Underneath, there are two input fields: "Customer VLAN ID (1-4094)" and "Service VLAN ID (1-4093)". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

To show the mapping table:

1. Click VLAN, Tunnel.
2. Select Configure Service from the Step list.
3. Select Show from the Action list.
4. Select an interface from the Port list.

Figure 77: Showing CVLAN to SPVLAN Mapping Entries

The screenshot shows the 'VLAN > Tunnel' configuration page. At the top, there are dropdown menus for 'Step: 2. Configure Service' and 'Action: Show'. Below that is a 'Port' dropdown set to '1'. The main content is a table titled 'Tunnel Service Subscriptions List' with a 'Total: 2' indicator. The table has two columns: 'Customer VLAN ID' and 'Service VLAN ID'. There are two rows of data: the first row has '1' in both columns, and the second row has '2' in the first column and '200' in the second column. At the bottom of the table are 'Delete' and 'Revert' buttons.

	Customer VLAN ID	Service VLAN ID
<input type="checkbox"/>	1	1
<input type="checkbox"/>	2	200

The preceding example sets the SVID to 99 in the outer tag for egress packets exiting port 1 when the packet’s CVID is 2. For a more detailed example, see the “switchport dot1q-tunnel service match cvid” command in the *CLI Reference Guide*.

Adding an Interface to a QinQ Tunnel

Follow the guidelines in the preceding section to set up a QinQ tunnel on the switch. Then use the VLAN > Tunnel (Configure Interface) page to set the tunnel mode for any participating interface.

Command Usage

- ◆ Use the Configure Global page to set the switch to QinQ mode before configuring a tunnel access port or tunnel uplink port (see [“Enabling QinQ Tunneling on the Switch” on page 158](#)). Also set the Tag Protocol Identifier (TPID) value of the tunnel access port if the attached client is using a nonstandard 2-byte ethertype to identify 802.1Q tagged frames.
- ◆ Then use the Configure Interface page to set the access interface on the edge switch to Access mode, and set the uplink interface on the switch attached to the service provider network to Uplink mode.

Parameters

These parameters are displayed:

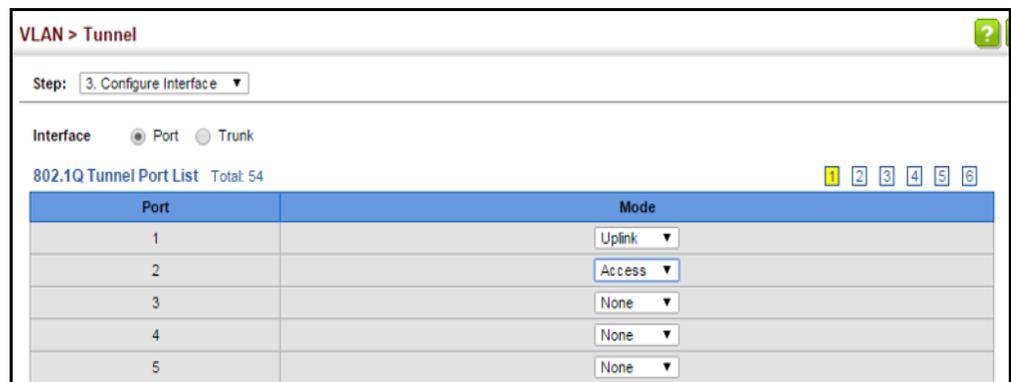
- ◆ **Interface** – Displays a list of ports or trunks.
 - **Port** – Port Identifier. (Range: 1-52)
 - **Trunk** – Trunk Identifier. (Range: 1-26)
- ◆ **Mode** – Sets the VLAN membership mode of the port.
 - **None** – The port operates in its normal VLAN mode. (This is the default.)
 - **Access** – Configures QinQ tunneling for a client access port to segregate and preserve customer VLAN IDs for traffic crossing the service provider network.
 - **Uplink** – Configures QinQ tunneling for an uplink port to another device within the service provider network.

Web Interface

To add an interface to a QinQ tunnel:

1. Click VLAN, Tunnel.
2. Select Configure Interface from the Step list.
3. Set the mode for any tunnel access port to Access and the tunnel uplink port to Uplink.
4. Click Apply.

Figure 78: Adding an Interface to a QinQ Tunnel



6

Address Table Settings

Switches store the addresses for all known devices. This information is used to pass traffic directly between the inbound and outbound ports. All the addresses learned by monitoring traffic are stored in the dynamic address table. You can also manually configure static addresses that are bound to a specific port.

This chapter describes the following topics:

- ◆ [MAC Address Learning](#) – Enables or disables address learning on an interface.
- ◆ [Static MAC Addresses](#) – Configures static entries in the address table.
- ◆ [Address Aging Time](#) – Sets timeout for dynamically learned entries.
- ◆ [Dynamic Address Cache](#) – Shows dynamic entries in the address table.
- ◆ [MAC Notification Traps](#) – Issue trap when a dynamic MAC address is added or removed.

Configuring MAC Address Learning

Use the [MAC Address > Learning Status](#) page to enable or disable MAC address learning on an interface.

Command Usage

- ◆ When MAC address learning is disabled, the switch immediately stops learning new MAC addresses on the specified interface. Only incoming traffic with source addresses stored in the static address table (see [“Setting Static Addresses” on page 165](#)) will be accepted as authorized to access the network through that interface.
- ◆ Dynamic addresses stored in the address table when MAC address learning is disabled are flushed from the system, and no dynamic addresses are subsequently learned until MAC address learning has been re-enabled. Any device not listed in the static address table that attempts to use the interface after MAC learning has been disabled will be prevented from accessing the switch.

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Port** – Port Identifier. (Range: 1-32/54)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-27)
- ◆ **Status** – The status of MAC address learning. (Default: Enabled)

Web Interface

To enable or disable MAC address learning:

1. Click MAC Address, Learning Status.
2. Set the learning status for any interface.
3. Click Apply.

Figure 79: Configuring MAC Address Learning

MAC Address > Learning Status

Interface Port Trunk

Port Learning Status List Total: 54

Port	Status
1	<input checked="" type="checkbox"/> Enabled
2	<input checked="" type="checkbox"/> Enabled
3	<input checked="" type="checkbox"/> Enabled
4	<input checked="" type="checkbox"/> Enabled
5	<input checked="" type="checkbox"/> Enabled

Setting Static Addresses

Use the MAC Address > Static page to configure static MAC addresses. A static address can be assigned to a specific interface on this switch. Static addresses are bound to the assigned interface and will not be moved. When a static address is seen on another interface, the address will be ignored and will not be written to the address table.

Command Usage

The static address for a host device can be assigned to a specific port within a specific VLAN. Use this command to add static addresses to the MAC Address Table. Static addresses have the following characteristics:

- ◆ Static addresses are bound to the assigned interface and will not be moved. When a static address is seen on another interface, the address will be ignored and will not be written to the address table.
- ◆ Static addresses will not be removed from the address table when a given interface link is down.
- ◆ A static address cannot be learned on another port until the address is removed from the table.

Parameters

These parameters are displayed:

Add Static Address

- ◆ **VLAN** – ID of configured VLAN. (Range: 1-4094)
- ◆ **Interface** – Port or trunk associated with the device assigned a static address.
- ◆ **MAC Address** – Physical address of a device mapped to this interface. Enter an address in the form of xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx.
- ◆ **Static Status** – Sets the time to retain the specified address.
 - Delete-on-reset - Assignment lasts until the switch is reset.
 - Permanent - Assignment is permanent. (This is the default.)

Show Static Address

The following additional fields are displayed on this web page:

Type – Displays the address configuration method. (Values: CPU, Config, or Security, the last of which indicates Port Security)

Life Time – The duration for which this entry applies. (Values: Delete On Reset, Delete On Timeout, Permanent)

Web Interface

To configure a static MAC address:

1. Click MAC Address, Static.
2. Select Add from the Action list.
3. Specify the VLAN, the port or trunk to which the address will be assigned, the MAC address, and the time to retain this entry.
4. Click Apply.

Figure 80: Configuring Static MAC Addresses

MAC Address > Static

Action: Add

VLAN: 1

Interface: Port 1 Trunk

MAC Address: 00-12-cf-94-34-da

Static Status: Permanent

Apply Revert

To show the static addresses in MAC address table:

1. Click MAC Address, Static.
2. Select Show from the Action list.

Figure 81: Displaying Static MAC Addresses

MAC Address > Static

Action: Show

Static MAC Address to Interface Mapping Table Total: 2

<input type="checkbox"/>	MAC Address	VLAN	Interface	Type	Life Time
<input type="checkbox"/>	00-00-0C-00-00-FD	1	CPU	CPU	Delete on Reset
<input type="checkbox"/>	00-12-CF-94-34-DA	1	Unit 1 / Port 1	Config	Permanent

Delete Revert

Changing the Aging Time

Use the MAC Address > Dynamic (Configure Aging) page to set the aging time for entries in the dynamic address table. The aging time is used to age out dynamically learned forwarding information.

Parameters

These parameters are displayed:

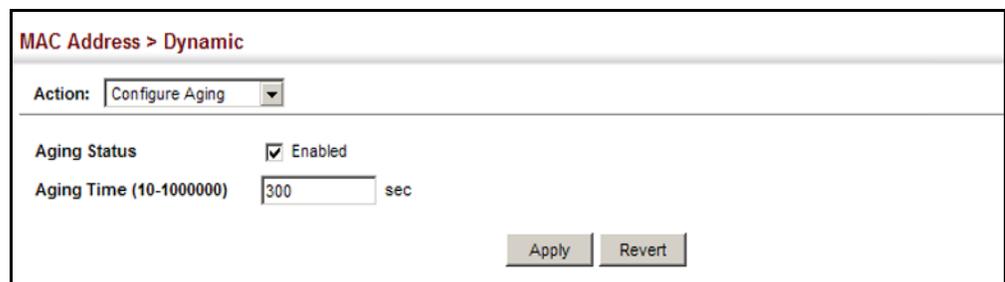
- ◆ **Aging Status** – Enables/disables the function.
- ◆ **Aging Time** – The time after which a learned entry is discarded. (Range: 10-1000000 seconds; Default: 300 seconds)

Web Interface

To set the aging time for entries in the dynamic address table:

1. Click MAC Address, Dynamic.
2. Select Configure Aging from the Action list.
3. Modify the aging status if required.
4. Specify a new aging time.
5. Click Apply.

Figure 82: Setting the Address Aging Time



The screenshot shows the 'MAC Address > Dynamic' configuration page. At the top, the breadcrumb 'MAC Address > Dynamic' is displayed. Below it, the 'Action:' dropdown menu is set to 'Configure Aging'. The 'Aging Status' is checked and labeled 'Enabled'. The 'Aging Time (10-1000000)' is set to '300' seconds. At the bottom right, there are 'Apply' and 'Revert' buttons.

Displaying the Dynamic Address Table

Use the MAC Address > Dynamic (Show Dynamic MAC) page to display the MAC addresses learned by monitoring the source address for traffic entering the switch. When the destination address for inbound traffic is found in the database, the packets intended for that address are forwarded directly to the associated port. Otherwise, the traffic is flooded to all ports.

Parameters

These parameters are displayed:

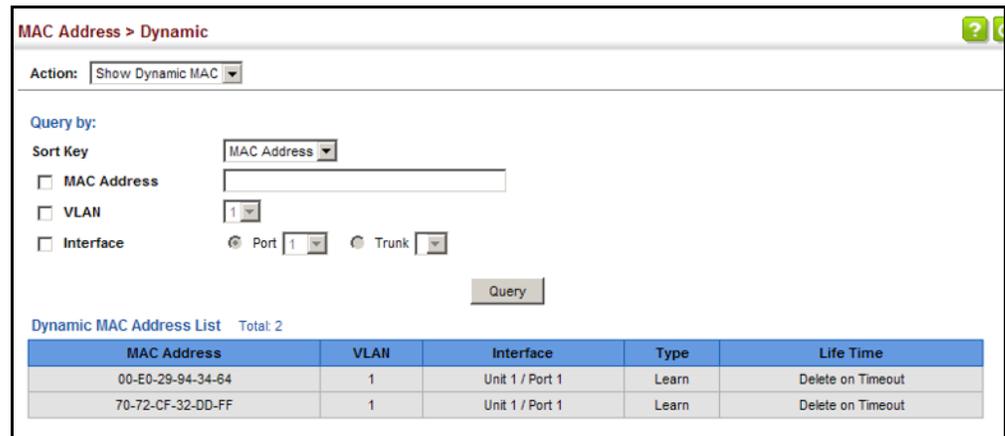
- ◆ **Sort Key** - You can sort the information displayed based on MAC address, VLAN or interface (port or trunk).
- ◆ **MAC Address** – Physical address associated with this interface.
- ◆ **VLAN** – ID of configured VLAN (1-4094).
- ◆ **Interface** – Indicates a port or trunk.
- ◆ **Type** – Shows that the entries in this table are learned.
- ◆ **Life Time** – Shows the time to retain the specified address.

Web Interface

To show the dynamic address table:

1. Click MAC Address, Dynamic.
2. Select Show Dynamic MAC from the Action list.
3. Select the Sort Key (MAC Address, VLAN, or Interface).
4. Enter the search parameters (MAC Address, VLAN, or Interface).
5. Click Query.

Figure 83: Displaying the Dynamic MAC Address Table



Clearing the Dynamic Address Table

Use the MAC Address > Dynamic (Clear Dynamic MAC) page to remove any learned entries from the forwarding database.

Parameters

These parameters are displayed:

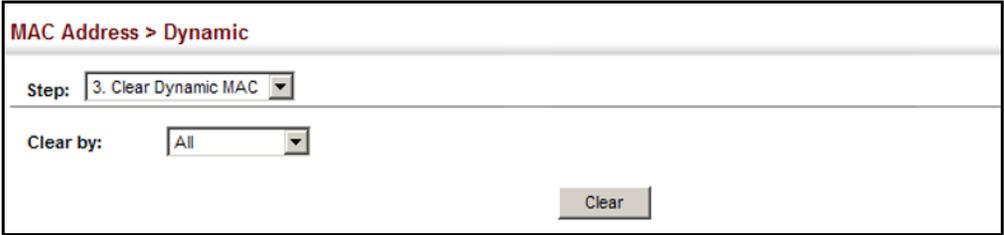
- ◆ **Clear by** – All entries can be cleared; or you can clear the entries for a specific MAC address, all the entries in a VLAN, or all the entries associated with a port or trunk.

Web Interface

To clear the entries in the dynamic address table:

1. Click MAC Address, Dynamic.
2. Select Clear Dynamic MAC from the Action list.
3. Select the method by which to clear the entries (i.e., All, MAC Address, VLAN, or Interface).
4. Enter information in the additional fields required for clearing entries by MAC Address, VLAN, or Interface.
5. Click Clear.

Figure 84: Clearing Entries in the Dynamic MAC Address Table



The screenshot shows a web interface for clearing dynamic MAC address entries. At the top, it says "MAC Address > Dynamic". Below that, there is a "Step:" dropdown menu currently set to "3. Clear Dynamic MAC". Underneath, there is a "Clear by:" dropdown menu set to "All". A "Clear" button is located at the bottom right of the form area.

Issuing MAC Address Traps

Use the MAC Address > MAC Notification pages to send SNMP traps (i.e., SNMP notifications) when a dynamic MAC address is added or removed.

Parameters

These parameters are displayed:

Configure Global

- ◆ **MAC Notification Traps** – Issues a trap when a dynamic MAC address is added or removed. (Default: Disabled)
- ◆ **MAC Notification Trap Interval** – Specifies the interval between issuing two consecutive traps. (Range: 1-3600 seconds; Default: 1 second)

Configure Interface

- ◆ **Interface** – Port or trunk identifier.
- ◆ **MAC Notification Trap** – Enables MAC authentication traps on the current interface. (Default: Disabled)
MAC authentication traps must be enabled at the global level for this attribute to take effect.

Web Interface

To enable MAC address traps at the global level:

1. Click MAC Address, MAC Notification.
2. Select Configure Global from the Step list.
3. Configure MAC notification traps and the transmission interval.
4. Click Apply.

Figure 85: Issuing MAC Address Traps (Global Configuration)

The screenshot shows the configuration page for MAC Address > MAC Notification. The 'Step' dropdown is set to '1. Configure Global'. There are two main settings: 'MAC Notification Traps' with an unchecked 'Enabled' checkbox, and 'MAC Notification Trap Interval (1-3600)' with a text input field containing the number '1' and the unit 'sec'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To enable MAC address traps at the interface level:

1. Click MAC Address, MAC Notification.
2. Select Configure Interface from the Step list.
3. Enable MAC notification traps for the required ports.
4. Click Apply.

Figure 86: Issuing MAC Address Traps (Interface Configuration)

The screenshot shows the configuration page for MAC Address > MAC Notification, Step 2: Configure Interface. The 'Step' dropdown is set to '2. Configure Interface'. Under the 'Interface' section, the 'Port' radio button is selected. Below this is a 'Port List' table with a total of 28 ports. The table has two columns: 'Port' and 'MAC Notification Trap'. The first five rows are visible, each showing a port number and an unchecked 'Enabled' checkbox. There are also three numbered tabs (1, 2, 3) at the top right of the table area.

Port	MAC Notification Trap
1	<input type="checkbox"/> Enabled
2	<input type="checkbox"/> Enabled
3	<input type="checkbox"/> Enabled
4	<input type="checkbox"/> Enabled
5	<input type="checkbox"/> Enabled

Spanning Tree Algorithm

This chapter describes the following basic topics:

- ◆ [Global Settings for STA](#) – Configures global bridge settings for STP, RSTP and MSTP.
- ◆ [Interface Settings for STA](#) – Configures interface settings for STA, including priority, path cost, link type, and designation as an edge port.
- ◆ [Global Settings for MSTP](#) – Sets the VLANs and associated priority assigned to an MST instance
- ◆ [Interface Settings for MSTP](#) – Configures interface settings for MSTP, including priority and path cost.

Overview

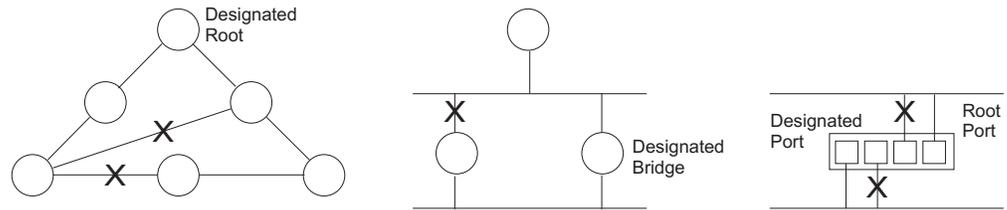
The Spanning Tree Algorithm (STA) can be used to detect and disable network loops, and to provide backup links between switches, bridges or routers. This allows the switch to interact with other bridging devices (that is, an STA-compliant switch, bridge or router) in your network to ensure that only one route exists between any two stations on the network, and provide backup links which automatically take over when a primary link goes down.

The spanning tree algorithms supported by this switch include these versions:

- ◆ STP – Spanning Tree Protocol (IEEE 802.1D)
- ◆ RSTP – Rapid Spanning Tree Protocol (IEEE 802.1w)
- ◆ MSTP – Multiple Spanning Tree Protocol (IEEE 802.1s)

STP – STP uses a distributed algorithm to select a bridging device (STP-compliant switch, bridge or router) that serves as the root of the spanning tree network. It selects a root port on each bridging device (except for the root device) which incurs the lowest path cost when forwarding a packet from that device to the root device. Then it selects a designated bridging device from each LAN which incurs the lowest path cost when forwarding a packet from that LAN to the root device. All ports connected to designated bridging devices are assigned as designated ports. After determining the lowest cost spanning tree, it enables all root ports and designated ports, and disables all other ports. Network packets are therefore only forwarded between root ports and designated ports, eliminating any possible network loops.

Figure 87: STP Root Ports and Designated Ports

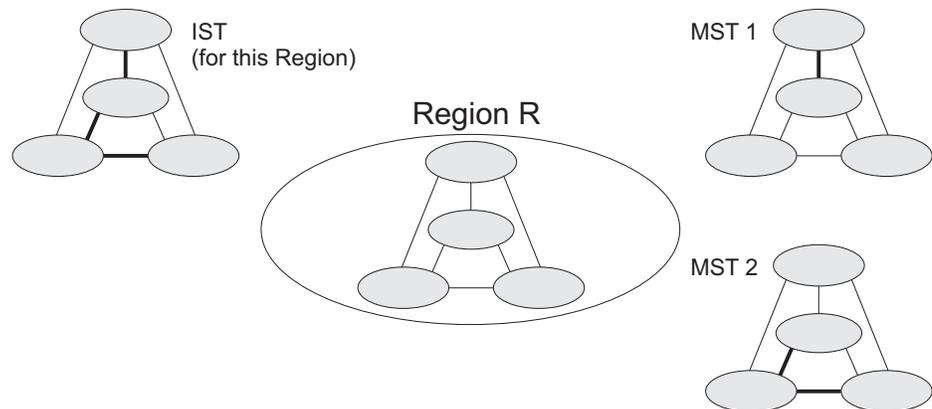


Once a stable network topology has been established, all bridges listen for Hello BPDUs (Bridge Protocol Data Units) transmitted from the Root Bridge. If a bridge does not get a Hello BPDU after a predefined interval (Maximum Age), the bridge assumes that the link to the Root Bridge is down. This bridge will then initiate negotiations with other bridges to reconfigure the network to reestablish a valid network topology.

RSTP – RSTP is designed as a general replacement for the slower, legacy STP. RSTP is also incorporated into MSTP. RSTP achieves much faster reconfiguration (i.e., around 1 to 3 seconds, compared to 30 seconds or more for STP) by reducing the number of state changes before active ports start learning, predefining an alternate route that can be used when a node or port fails, and retaining the forwarding database for ports insensitive to changes in the tree structure when reconfiguration occurs.

MSTP – When using STP or RSTP, it may be difficult to maintain a stable path between all VLAN members. Frequent changes in the tree structure can easily isolate some of the group members. MSTP (which is based on RSTP for fast convergence) is designed to support independent spanning trees based on VLAN groups. Using multiple spanning trees can provide multiple forwarding paths and enable load balancing. One or more VLANs can be grouped into a Multiple Spanning Tree Instance (MSTI). MSTP builds a separate Multiple Spanning Tree (MST) for each instance to maintain connectivity among each of the assigned VLAN groups. MSTP then builds an Internal Spanning Tree (IST) for the Region containing all commonly configured MSTP bridges.

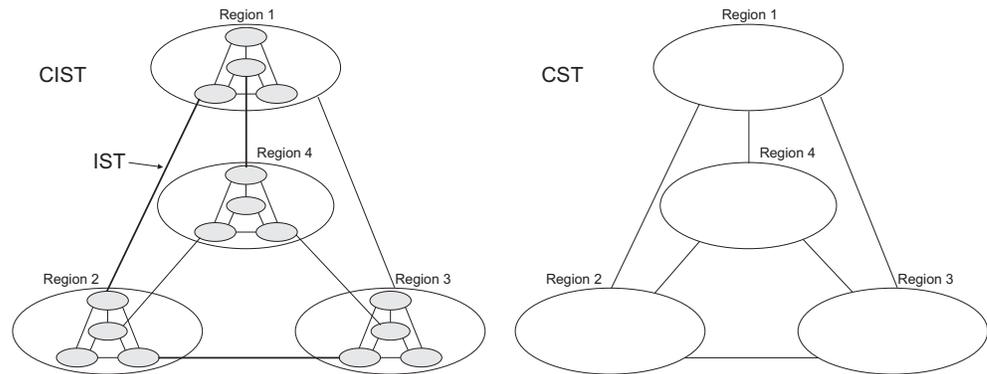
Figure 88: MSTP Region, Internal Spanning Tree, Multiple Spanning Tree



An MST Region consists of a group of interconnected bridges that have the same MST Configuration Identifiers (including the Region Name, Revision Level and

Configuration Digest – see “Configuring Multiple Spanning Trees” on page 188). An MST Region may contain multiple MSTP Instances. An Internal Spanning Tree (IST) is used to connect all the MSTP switches within an MST region. A Common Spanning Tree (CST) interconnects all adjacent MST Regions, and acts as a virtual bridge node for communications with STP or RSTP nodes in the global network.

Figure 89: Spanning Tree – Common Internal, Common, Internal



MSTP connects all bridges and LAN segments with a single Common and Internal Spanning Tree (CIST). The CIST is formed as a result of the running spanning tree algorithm between switches that support the STP, RSTP, MSTP protocols.

Once you specify the VLANs to include in a Multiple Spanning Tree Instance (MSTI), the protocol will automatically build an MSTI tree to maintain connectivity among each of the VLANs. MSTP maintains contact with the global network because each instance is treated as an RSTP node in the Common Spanning Tree (CST).

Configuring Global Settings for STA

Use the Spanning Tree > STA (Configure Global - Configure) page to configure global settings for the spanning tree that apply to the entire switch.

Command Usage

◆ Spanning Tree Protocol³

This option uses RSTP set to STP forced compatibility mode. It uses RSTP for the internal state machine, but sends only 802.1D BPDUs. This creates one spanning tree instance for the entire network. If multiple VLANs are implemented on a network, the path between specific VLAN members may be inadvertently disabled to prevent network loops, thus isolating group members. When operating multiple VLANs, we recommend selecting the MSTP option.

3. STP and RSTP BPDUs are transmitted as untagged frames, and will cross any VLAN boundaries.

◆ **Rapid Spanning Tree Protocol³**

RSTP supports connections to either STP or RSTP nodes by monitoring the incoming protocol messages and dynamically adjusting the type of protocol messages the RSTP node transmits, as described below:

- **STP Mode** – If the switch receives an 802.1D BPDU (i.e., STP BPDU) after a port's migration delay timer expires, the switch assumes it is connected to an 802.1D bridge and starts using only 802.1D BPDUs.
- **RSTP Mode** – If RSTP is using 802.1D BPDUs on a port and receives an RSTP BPDU after the migration delay expires, RSTP restarts the migration delay timer and begins using RSTP BPDUs on that port.

◆ **Multiple Spanning Tree Protocol**

MSTP generates a unique spanning tree for each instance. This provides multiple pathways across the network, thereby balancing the traffic load, preventing wide-scale disruption when a bridge node in a single instance fails, and allowing for faster convergence of a new topology for the failed instance.

- To allow multiple spanning trees to operate over the network, you must configure a related set of bridges with the same MSTP configuration, allowing them to participate in a specific set of spanning tree instances.
- A spanning tree instance can exist only on bridges that have compatible VLAN instance assignments.
- Be careful when switching between spanning tree modes. Changing modes stops all spanning-tree instances for the previous mode and restarts the system in the new mode, temporarily disrupting user traffic.

Parameters

These parameters are displayed:

Basic Configuration of Global Settings

- ◆ **Spanning Tree Status** – Enables/disables STA on this switch. (Default: Enabled)
- ◆ **Spanning Tree Type** – Specifies the type of spanning tree used on this switch:
 - **STP**: Spanning Tree Protocol (IEEE 802.1D); i.e., when this option is selected, the switch will use RSTP set to STP forced compatibility mode).
 - **RSTP**: Rapid Spanning Tree (IEEE 802.1w); RSTP is the default.
 - **MSTP**: Multiple Spanning Tree (IEEE 802.1s)
- ◆ **Priority** – Bridge priority is used in selecting the root device, root port, and designated port. The device with the highest priority becomes the STA root device. However, if all devices have the same priority, the device with the

lowest MAC address will then become the root device. (Note that lower numeric values indicate higher priority.)

- Default: 32768
- Range: 0-61440, in steps of 4096
- Options: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, 61440

Advanced Configuration Settings

The following attributes are based on RSTP, but also apply to STP since the switch uses a backwards-compatible subset of RSTP to implement STP, and also apply to MSTP which is based on RSTP according to the standard:

- ◆ **Path Cost Method** – The path cost is used to determine the best path between devices. The path cost method is used to determine the range of values that can be assigned to each interface.
 - Long: Specifies 32-bit based values that range from 1-200,000,000. (This is the default.)
 - Short: Specifies 16-bit based values that range from 1-65535.
- ◆ **Transmission Limit** – The maximum transmission rate for BPDUs is specified by setting the minimum interval between the transmission of consecutive protocol messages. (Range: 1-10; Default: 3)

When the Switch Becomes Root

- ◆ **Hello Time** – Interval (in seconds) at which the root device transmits a configuration message.
 - Default: 2
 - Minimum: 1
 - Maximum: The lower of 10 or $[(\text{Max. Message Age} / 2) - 1]$
- ◆ **Maximum Age** – The maximum time (in seconds) a device can wait without receiving a configuration message before attempting to reconverge. All device ports (except for designated ports) should receive configuration messages at regular intervals. Any port that ages out STA information (provided in the last configuration message) becomes the designated port for the attached LAN. If it is a root port, a new root port is selected from among the device ports attached to the network. (References to “ports” in this section mean “interfaces,” which includes both ports and trunks.)
 - Default: 20
 - Minimum: The higher of 6 or $[2 \times (\text{Hello Time} + 1)]$
 - Maximum: The lower of 40 or $[2 \times (\text{Forward Delay} - 1)]$

- ◆ **Forward Delay** – The maximum time (in seconds) this device will wait before changing states (i.e., discarding to learning to forwarding). This delay is required because every device must receive information about topology changes before it starts to forward frames. In addition, each port needs time to listen for conflicting information that would make it return to a discarding state; otherwise, temporary data loops might result.
 - Default: 15
 - Minimum: The higher of 4 or $[(\text{Max. Message Age} / 2) + 1]$
 - Maximum: 30

RSTP does not depend on the forward delay timer in most cases. It is able to confirm that a port can transition to the forwarding state without having to rely on any timer configuration. To achieve fast convergence, RSTP relies on the use of edge ports, and automatic detection of point-to-point link types, both of which allow a port to directly transition to the forwarding state.

Configuration Settings for MSTP

- ◆ **Max Instance Numbers** – The maximum number of MSTP instances to which this switch can be assigned.
- ◆ **Configuration Digest** – An MD5 signature key that contains the VLAN ID to MST ID mapping table. In other words, this key is a mapping of all VLANs to the CIST.
- ◆ **Region Revision**⁴ – The revision for this MSTI. (Range: 0-65535; Default: 0)
- ◆ **Region Name**⁴ – The name for this MSTI. (Maximum length: 32 characters; switch's MAC address)
- ◆ **Max Hop Count** – The maximum number of hops allowed in the MST region before a BPDU is discarded. (Range: 1-40; Default: 20)

Web Interface

To configure global STA settings:

1. Click Spanning Tree, STA.
2. Select Configure Global from the Step list.
3. Select Configure from the Action list.
4. Modify any of the required attributes. Note that the parameters displayed for the spanning tree types (STP, RSTP, MSTP) varies as described in the preceding section.
5. Click Apply

4. The MST name and revision number are both required to uniquely identify an MST region.

Figure 90: Configuring Global Settings for STA (STP)

Spanning Tree > STA

Step: 1. Configure Global Action: Configure

Spanning Tree Status Enabled

Spanning Tree Type **STP**

Priority (0-61440, in steps of 4096) 32768

Advanced:

Path Cost Method Long

Transmission Limit (1-10) 3

When the Switch Becomes Root:

Hello Time (1-10) 2 sec

Maximum Age (6-40) 20 sec

Forward Delay (4-30) 15 sec

Note: $2 * (\text{Hello Time} + 1) \leq \text{Max Age} \leq 2 * (\text{Forward Delay} - 1)$

Apply Revert

Figure 91: Configuring Global Settings for STA (RSTP)

Spanning Tree > STA

Step: 1. Configure Global Action: Configure

Spanning Tree Status Enabled

Spanning Tree Type **RSTP**

Priority (0-61440, in steps of 4096) 32768

Advanced:

Path Cost Method Long

Transmission Limit (1-10) 3

When the Switch Becomes Root:

Hello Time (1-10) 2 sec

Maximum Age (6-40) 20 sec

Forward Delay (4-30) 15 sec

Note: $2 * (\text{Hello Time} + 1) \leq \text{Max Age} \leq 2 * (\text{Forward Delay} - 1)$

Apply Revert

Figure 92: Configuring Global Settings for STA (MSTP)

The screenshot shows the 'Spanning Tree > STA' configuration page. At the top, there is a breadcrumb 'Spanning Tree > STA'. Below it, a 'Step' dropdown is set to '1. Configure Global' and an 'Action' dropdown is set to 'Configure'. The main configuration area is divided into several sections:

- Spanning Tree Status:** A checkbox labeled 'Enabled' is checked.
- Spanning Tree Type:** A dropdown menu is set to 'MSTP'.
- Priority (0-61440, in steps of 4096):** A text input field contains the value '32768'.
- Advanced:**
 - Path Cost Method:** A dropdown menu is set to 'Long'.
 - Transmission Limit (1-10):** A text input field contains the value '3'.
- When the Switch Becomes Root:**
 - Hello Time (1-10):** A text input field contains '2' with 'sec' to its right.
 - Maximum Age (6-40):** A text input field contains '20' with 'sec' to its right.
 - Forward Delay (4-30):** A text input field contains '15' with 'sec' to its right.
- Note:** $2 * (\text{Hello Time} + 1) \leq \text{Max Age} \leq 2 * (\text{Forward Delay} - 1)$
- MSTP Configuration:**
 - Max Instance Numbers:** A text input field contains '33'.
 - Configuration Digest:** A text input field contains '0xAC36177F50283CD4B883821D8AB26DE62'.
 - Region Revision (0-65535):** A text input field contains '0'.
 - Region Name:** A text input field contains '00 00 e8 93 82 a0'.
 - Max Hop Count (1-40):** A text input field contains '20'.

At the bottom right of the configuration area, there are two buttons: 'Apply' and 'Revert'.

Displaying Global Settings for STA

Use the Spanning Tree > STA (Configure Global - Show Information) page to display a summary of the current bridge STA information that applies to the entire switch.

Parameters

The parameters displayed are described in the preceding section, except for the following items:

- ◆ **Bridge ID** – A unique identifier for this bridge, consisting of the bridge priority, the MST Instance ID 0 for the Common Spanning Tree when spanning tree type is set to MSTP, and MAC address (where the address is taken from the switch system).
- ◆ **Designated Root** – The priority and MAC address of the device in the Spanning Tree that this switch has accepted as the root device.
- ◆ **Root Port** – The number of the port on this switch that is closest to the root. This switch communicates with the root device through this port. If there is no root port, then this switch has been accepted as the root device of the Spanning Tree network.

- ◆ **Root Path Cost** – The path cost from the root port on this switch to the root device.
- ◆ **Topology Changes** – The number of times the Spanning Tree has been reconfigured.
- ◆ **Last Topology Change** – Time since the Spanning Tree was last reconfigured.

Web Interface

To display global STA settings:

1. Click Spanning Tree, STA.
2. Select Configure Global from the Step list.
3. Select Show Information from the Action list.

Figure 93: Displaying Global Settings for STA

The screenshot shows the 'Spanning Tree > STA' configuration page. At the top, there are two dropdown menus: 'Step: 1. Configure Global' and 'Action: Show Information'. Below these is a section titled 'Spanning Tree Information' containing a table of settings.

Spanning Tree Information			
Spanning Tree Status	Enabled	Spanning Tree Type	RSTP
Designated Root	32768.7072CF800E50	Bridge ID	32768.7072CF800E50
Root Port	0	Max Age	20 sec
Root Path Cost	0	Hello Time	2 sec
Topology Changes	0	Forward Delay	15 sec
Last Topology Change	0 days, 2 hours, 13 minutes, 3 seconds		

Configuring Interface Settings for STA

Use the Spanning Tree > STA (Configure Interface - Configure) page to configure RSTP and MSTP attributes for specific interfaces, including port priority, path cost, link type, and edge port. You may use a different priority or path cost for ports of the same media type to indicate the preferred path, link type to indicate a point-to-point connection or shared-media connection, and edge port to indicate if the attached device can support fast forwarding. (References to “ports” in this section means “interfaces,” which includes both ports and trunks.)

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Spanning Tree** – Enables/disables STA on this interface. (Default: Enabled)

- ◆ **BPDU Flooding** – Enables/disables the flooding of BPDUs to other ports when global spanning tree is disabled ([page 175](#)) or when spanning tree is disabled on a specific port. When flooding is enabled, BPDUs are flooded to all other ports on the switch or to all other ports within the receiving port's native VLAN as specified by the Spanning Tree BPDU Flooding attribute ([page 175](#)).
- ◆ **Priority** – Defines the priority used for this port in the Spanning Tree Protocol. If the path cost for all ports on a switch are the same, the port with the highest priority (i.e., lowest value) will be configured as an active link in the Spanning Tree. This makes a port with higher priority less likely to be blocked if the Spanning Tree Protocol is detecting network loops. Where more than one port is assigned the highest priority, the port with lowest numeric identifier will be enabled.
 - Default: 128
 - Range: 0-240, in steps of 16
- ◆ **Admin Path Cost** – This parameter is used by the STA to determine the best path between devices. Therefore, lower values should be assigned to ports attached to faster media, and higher values assigned to ports with slower media. Note that path cost takes precedence over port priority. (Range: 0 for auto-configuration, 1-65535 for the short path cost method⁵, 1-200,000,000 for the long path cost method)

By default, the system automatically detects the speed and duplex mode used on each port, and configures the path cost according to the values shown below. Path cost "0" is used to indicate auto-configuration mode. When the short path cost method is selected and the default path cost recommended by the IEEE 802.1w standard exceeds 65,535, the default is set to 65,535.

Table 11: Recommended STA Path Cost Range

Port Type	IEEE 802.1D-1998	IEEE 802.1w-2001
Ethernet	50-600	200,000-20,000,000
Fast Ethernet	10-60	20,000-2,000,000
Gigabit Ethernet	3-10	2,000-200,000
10G Ethernet	1-5	200-20,000
40G Ethernet	1-65535 ¹	20-2,000 ¹

¹ Undefined in standard.

5. Refer to [“Configuring Global Settings for STA” on page 175](#) for information on setting the path cost method. The range displayed on the STA interface configuration page shows the maximum value for path cost. However, note that the switch still enforces the rules for path cost based on the specified path cost method (long or short).

Table 12: Default STA Path Costs

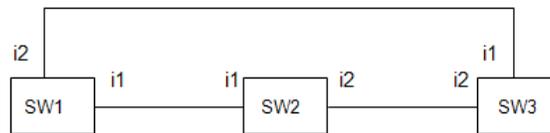
Port Type	Short Path Cost (IEEE 802.1D-1998)	Long Path Cost (IEEE 802.1D-2004)
Ethernet	65,535	1,000,000
Fast Ethernet	65,535	100,000
Gigabit Ethernet	10,000	10,000
10G Ethernet	1,000	1,000
40G Ethernet	65535 ¹	2,000,000 ²

1 Undefined in standard, but recommended setting is 250.

2 Code does not support 40G path cost, and therefore defaults to 10M half duplex cost.

Administrative path cost cannot be used to directly determine the root port on a switch. Connections to other devices use IEEE 802.1Q-2005 to determine the root port as in the following example.

Figure 94: Determining the Root Port



For BPDU messages received by i1 on SW3, the path cost is 0.

For BPDU messages received by i2 on SW3, the path cost is that of i1 on SW2.

The root path cost for i1 on SW3 used to compete for the role of root port is 0 + path cost of i1 on SW3; 0 since i1 is directly connected to the root bridge.

If the path cost of i1 on SW2 is never configured/changed, it is 10000.

Then the root path cost for i2 on SW3 used to compete for the role of root port is 10000 + path cost of i2 on SW3.

The path cost of i1 on SW3 is also 10000 if not configured/changed.

Then even if the path cost of i2 on SW3 is configured/changed to 0, these ports will still have the same root path cost, and it will be impossible for i2 to become the root port just by changing its path cost on SW3.

For RSTP mode, the root port can be determined simply by adjusting the path cost of i1 on SW2. However, for MSTP mode, it is impossible to achieve this only by changing the path cost because external path cost is not added in the same region, and the regional root for i1 is SW1, but for i2 is SW2.

- ◆ **Admin Link Type** – The link type attached to this interface.
 - Point-to-Point – A connection to exactly one other bridge.
 - Shared – A connection to two or more bridges.
 - Auto – The switch automatically determines if the interface is attached to a point-to-point link or to shared media. (This is the default setting.)

- ◆ **Root Guard** – STA allows a bridge with a lower bridge identifier (or same identifier and lower MAC address) to take over as the root bridge at any time. Root Guard can be used to ensure that the root bridge is not formed at a suboptimal location. Root Guard should be enabled on any designated port connected to low-speed bridges which could potentially overload a slower link by taking over as the root port and forming a new spanning tree topology. It could also be used to form a border around part of the network where the root bridge is allowed. (Default: Disabled)

- ◆ **Admin Edge Port** – Since end nodes **cannot** cause forwarding loops, they can pass directly through to the spanning tree forwarding state. Specifying Edge Ports provides quicker convergence for devices such as workstations or servers, retains the current forwarding database to reduce the amount of frame flooding required to rebuild address tables during reconfiguration events, does not cause the spanning tree to initiate reconfiguration when the interface changes state, and also overcomes other STA-related timeout problems. However, remember that Edge Port should only be enabled for ports connected to an end-node device. (Default: Auto)
 - **Enabled** – Manually configures a port as an Edge Port.
 - **Disabled** – Disables the Edge Port setting.
 - **Auto** – The port will be automatically configured as an edge port if the edge delay time expires without receiving any RSTP or MSTP BPDUs. Note that edge delay time (802.1D-2004 17.20.4) equals the protocol migration time if a port's link type is point-to-point (which is 3 seconds as defined in IEEE 802.3D-2004 17.20.4); otherwise it equals the spanning tree's maximum age for configuration messages (see maximum age under ["Configuring Global Settings for STA" on page 175](#)).

An interface cannot function as an edge port under the following conditions:

- If spanning tree mode is set to STP ([page 175](#)), edge-port mode cannot automatically transition to operational edge-port state using the automatic setting.
- If an interface is in forwarding state and its role changes, the interface cannot continue to function as an edge port even if the edge delay time has expired.
- If the port does not receive any BPDUs after the edge delay timer expires, its role changes to designated port and it immediately enters forwarding state (see ["Displaying Interface Settings for STA" on page 186](#)).

When edge port is set as auto, the operational state is determined automatically by the Bridge Detection State Machine described in 802.1D-2004, where the edge port state may change dynamically based on environment changes (e.g., receiving a BPDU or not within the required interval).

- ◆ **BPDU Guard** – This feature protects edge ports from receiving BPDUs. It prevents loops by shutting down an edge port when a BPDU is received instead of putting it into the spanning tree discarding state. In a valid configuration, configured edge ports should not receive BPDUs. If an edge port receives a BPDU an invalid configuration exists, such as a connection to an unauthorized device. The BPDU guard feature provides a secure response to invalid configurations because an administrator must manually enable the port. (Default: Disabled)

BPDU guard can only be configured on an interface if the edge port attribute is not disabled (that is, if the edge port is set to enabled or auto).

- ◆ **BPDU Guard Auto Recovery** – Automatically re-enables an interface after the specified interval. (Default: Disabled)

- ◆ **BPDU Guard Auto Recovery Interval** – The time to wait before re-enabling an interface. (Range: 30-86400 seconds; Default: 300 seconds)

- ◆ **BPDU Filter** – BPDU filtering allows you to avoid transmitting BPDUs on configured edge ports that are connected to end nodes. By default, STA sends BPDUs to all ports regardless of whether administrative edge is enabled on a port. BPDU filtering is configured on a per-port basis. (Default: Disabled)

BPDU filter can only be configured on an interface if the edge port attribute is not disabled (that is, if edge port is set to enabled or auto).

- ◆ **Migration** – If at any time the switch detects STP BPDUs, including Configuration or Topology Change Notification BPDUs, it will automatically set the selected interface to forced STP-compatible mode. However, you can also use the Protocol Migration button to manually re-check the appropriate BPDU format (RSTP or STP-compatible) to send on the selected interfaces. (Default: Disabled)

Web Interface

To configure interface settings for STA:

1. Click Spanning Tree, STA.
2. Select Configure Interface from the Step list.
3. Select Configure from the Action list.
4. Modify any of the required attributes.
5. Click Apply.

Figure 95: Configuring Interface Settings for STA

The screenshot shows the 'Spanning Tree > STA' configuration page. At the top, there are dropdowns for 'Step: 2. Configure Interface' and 'Action: Configure'. Below that, there are radio buttons for 'Interface' set to 'Port'. A 'Port List' section shows 'Total: 54' ports. A table below displays settings for five ports. The table has columns for Port, Spanning Tree, Priority, Admin Path Cost, Admin Link Type, Root Guard, Admin Edge Port, BPDU Guard, BPDU Guard Auto Recovery, BPDU Guard Auto Recovery Interval, BPDU Filter, and Migration. All 'Spanning Tree' checkboxes are checked. 'Priority' is set to 128, 'Admin Path Cost' to 0, and 'Admin Link Type' to 'Auto'. 'Root Guard' is checked. 'Admin Edge Port' is set to 'Auto'. 'BPDU Guard' and 'BPDU Guard Auto Recovery' are checked. 'BPDU Guard Auto Recovery Interval' is set to 300. 'BPDU Filter' and 'Migration' are unchecked.

Port	Spanning Tree	Priority (0-240, in steps of 16)	Admin Path Cost (0-200000000, 0: Auto)	Admin Link Type	Root Guard	Admin Edge Port	BPDU Guard	BPDU Guard Auto Recovery	BPDU Guard Auto Recovery Interval (30-86400)	BPDU Filter	Migration
1	<input checked="" type="checkbox"/> Enabled	128	0	Auto	<input checked="" type="checkbox"/> Enabled	Auto	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	300	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
2	<input checked="" type="checkbox"/> Enabled	128	0	Auto	<input checked="" type="checkbox"/> Enabled	Auto	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	300	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
3	<input checked="" type="checkbox"/> Enabled	128	0	Auto	<input checked="" type="checkbox"/> Enabled	Auto	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	300	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
4	<input checked="" type="checkbox"/> Enabled	128	0	Auto	<input checked="" type="checkbox"/> Enabled	Auto	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	300	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
5	<input checked="" type="checkbox"/> Enabled	128	0	Auto	<input checked="" type="checkbox"/> Enabled	Auto	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	300	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled

Displaying Interface Settings for STA

Use the Spanning Tree > STA (Configure Interface - Show Information) page to display the current status of ports or trunks in the Spanning Tree.

Parameters

These parameters are displayed:

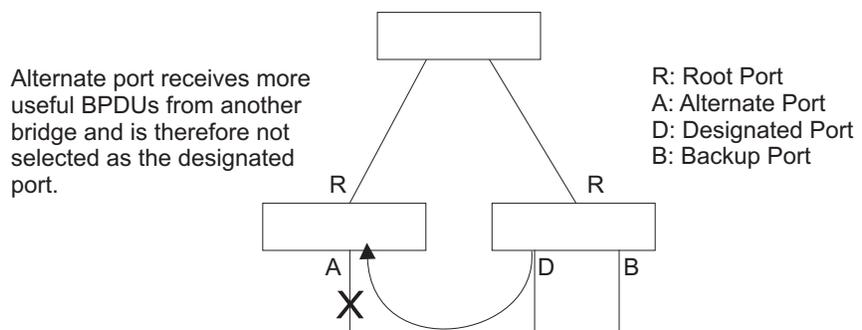
- ◆ **Spanning Tree** – Shows if STA has been enabled on this interface.
- ◆ **STA Status** – Displays current state of this port within the Spanning Tree:
 - **Discarding** - Port receives STA configuration messages, but does not forward packets.
 - **Learning** - Port has transmitted configuration messages for an interval set by the Forward Delay parameter without receiving contradictory information. Port address table is cleared, and the port begins learning addresses.
 - **Forwarding** - Port forwards packets, and continues learning addresses.

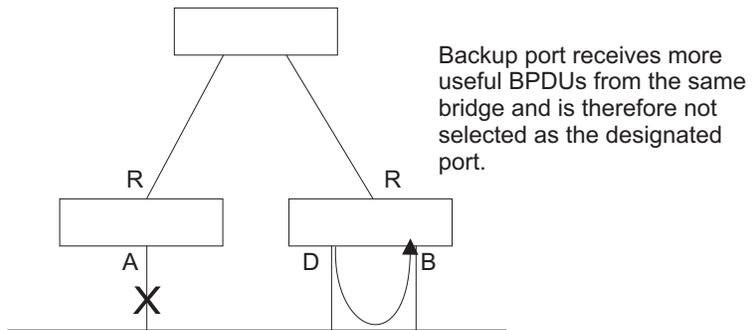
The rules defining port status are:

- A port on a network segment with no other STA compliant bridging device is always forwarding.
- If two ports of a switch are connected to the same segment and there is no other STA device attached to this segment, the port with the smaller ID forwards packets and the other is discarding.

- All ports are discarding when the switch is booted, then some of them change state to learning, and then to forwarding.
- ◆ **Forward Transitions** – The number of times this port has transitioned from the Learning state to the Forwarding state.
- ◆ **Designated Cost** – The cost for a packet to travel from this port to the root in the current Spanning Tree configuration. The slower the media, the higher the cost.
- ◆ **Designated Bridge** – The bridge priority and MAC address of the device through which this port must communicate to reach the root of the Spanning Tree.
- ◆ **Designated Port** – The port priority and number of the port on the designated bridging device through which this switch must communicate with the root of the Spanning Tree.
- ◆ **Oper Path Cost** – The contribution of this port to the path cost of paths towards the spanning tree root which include this port.
- ◆ **Oper Link Type** – The operational point-to-point status of the LAN segment attached to this interface. This parameter is determined by manual configuration or by auto-detection, as described for Admin Link Type in STA Port Configuration on [page 181](#).
- ◆ **Oper Edge Port** – This parameter is initialized to the setting for Admin Edge Port in STA Port Configuration on [page 181](#) (i.e., true or false), but will be set to false if a BPDU is received, indicating that another bridge is attached to this port.
- ◆ **Port Role** – Roles are assigned according to whether the port is part of the active topology, that is the best port connecting a non-root bridge to the root bridge (i.e., **root** port), connecting a LAN through the bridge to the root bridge (i.e., **designated** port), is the MSTI regional root (i.e., **master** port), or is an **alternate** or **backup** port that may provide connectivity if other bridges, bridge ports, or LANs fail or are removed. The role is set to disabled (i.e., **disabled** port) if a port has no role within the spanning tree.

Figure 96: STA Port Roles





The criteria used for determining the port role is based on root bridge ID, root path cost, designated bridge, designated port, port priority, and port number, in that order and as applicable to the role under question.

Web Interface

To display interface settings for STA:

1. Click Spanning Tree, STA.
2. Select Configure Interface from the Step list.
3. Select Show Information from the Action list.

Figure 97: Displaying Interface Settings for STA

Spanning Tree > STA										
Step:		2. Configure Interface		Action:		Show Information				
Interface <input checked="" type="radio"/> Port <input type="radio"/> Trunk										
Spanning Tree Port List Total: 54										
Port	Spanning Tree	STA Status	Forward Transitions	Designated Cost	Designated Bridge	Designated Port	Oper Path Cost	Oper Link Type	Oper Edge Port	Port Role
1	Enabled	Discarding	0	250	32768.7072CFEA1B71	128.1	1000	Point-to-Point	Disabled	Disabled
2	Enabled	Discarding	0	250	32768.7072CFEA1B71	128.2	1000	Point-to-Point	Disabled	Disabled
3	Enabled	Discarding	0	250	32768.7072CFEA1B71	128.3	1000	Point-to-Point	Disabled	Disabled
4	Enabled	Discarding	0	250	32768.7072CFEA1B71	128.4	1000	Point-to-Point	Disabled	Disabled
5	Enabled	Discarding	0	250	32768.7072CFEA1B71	128.5	1000	Point-to-Point	Disabled	Disabled

Configuring Multiple Spanning Trees

Use the Spanning Tree > MSTP (Configure Global) page to create an MSTP instance, or to add VLAN groups to an MSTP instance.

Command Usage

MSTP generates a unique spanning tree for each instance. This provides multiple pathways across the network, thereby balancing the traffic load, preventing wide-scale disruption when a bridge node in a single instance fails, and allowing for faster convergence of a new topology for the failed instance.

By default all VLANs are assigned to the Internal Spanning Tree (MST Instance 0) that connects all bridges and LANs within the MST region. This switch supports up to 33 instances. You should try to group VLANs which cover the same general area of your network. However, remember that you must configure all bridges within the same MSTI Region (page 175) with the same set of instances, and the same instance (on each bridge) with the same set of VLANs. Also, note that RSTP treats each MSTI region as a single node, connecting all regions to the Common Spanning Tree.

To use multiple spanning trees:

1. Set the spanning tree type to MSTP (page 175).
2. Enter the spanning tree priority for the selected MST instance on the Spanning Tree > MSTP (Configure Global - Add) page.
3. Add the VLANs that will share this MSTI on the Spanning Tree > MSTP (Configure Global - Add Member) page.



Note: All VLANs are automatically added to the IST (Instance 0).

To ensure that the MSTI maintains connectivity across the network, you must configure a related set of bridges with the same MSTI settings.

Parameters

These parameters are displayed:

- ◆ **MST ID** – Instance identifier to configure. (Range: 0-4094)
- ◆ **VLAN ID** – VLAN to assign to this MST instance. (Range: 1-4094)
- ◆ **Priority** – The priority of a spanning tree instance. (Range: 0-61440 in steps of 4096; Options: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, 61440; Default: 32768)

Web Interface

To create instances for MSTP:

1. Click Spanning Tree, MSTP.
2. Select Configure Global from the Step list.
3. Select Add from the Action list.
4. Specify the MST instance identifier and the initial VLAN member. Additional member can be added using the Spanning Tree > MSTP (Configure Global - Add Member) page. If the priority is not specified, the default value 32768 is used.

5. Click Apply.

Figure 98: Creating an MST Instance

The screenshot shows the 'Spanning Tree > MSTP' configuration page. At the top, there is a breadcrumb 'Spanning Tree > MSTP'. Below it, there are two dropdown menus: 'Step: 1. Configure Global' and 'Action: Add'. The main configuration area contains three input fields: 'MST ID (0-4094)' with the value '1', 'VLAN ID (1-4094)' with the value '1', and 'Priority (0-61440, in steps of 4096)' which is empty. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

To show the MSTP instances:

1. Click Spanning Tree, MSTP.
2. Select Configure Global from the Step list.
3. Select Show from the Action list.

Figure 99: Displaying MST Instances

The screenshot shows the 'Spanning Tree > MSTP' configuration page with the 'Action' dropdown set to 'Show'. Below the configuration fields, there is a table titled 'MST List Total: 2'. The table has a header row with a checkbox and 'MST ID'. There are two data rows: one with MST ID 0 and one with MST ID 1. At the bottom right, there are two buttons: 'Delete' and 'Revert'.

<input type="checkbox"/>	MST ID
<input type="checkbox"/>	0
<input type="checkbox"/>	1

To modify the priority for an MST instance:

1. Click Spanning Tree, MSTP.
2. Select Configure Global from the Step list.
3. Select Modify from the Action list.
4. Modify the priority for an MSTP Instance.
5. Click Apply.

Figure 100: Modifying the Priority for an MST Instance

Spanning Tree > MSTP

Step: 1. Configure Global Action: Modify

MST Details List Total: 2

MST ID	Priority (0-61440, in steps of 4096)
0	0
1	32768

Apply Revert

To display global settings for MSTP:

1. Click Spanning Tree, MSTP.
2. Select Configure Global from the Step list.
3. Select Show Information from the Action list.
4. Select an MST ID. The attributes displayed on this page are described under “Displaying Global Settings for STA” on page 180.

Figure 101: Displaying Global Settings for an MST Instance

Spanning Tree > MSTP

Step: 1. Configure Global Action: Show Information

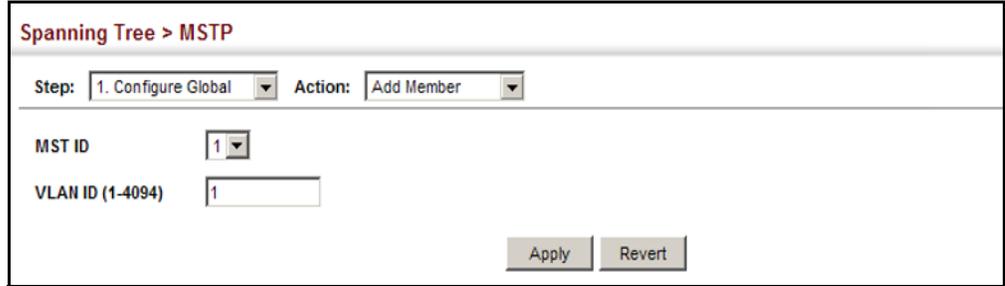
MST ID 1

Priority	0	Designated Root	32768.0030F1245660
Bridge ID	20	Root Port	2
Max Age	15 sec	Root Path Cost	32768.000001010010
Hello Time	23 sec	Configuration Changes	500000
Forward Delay	2 sec	Last Topology Change	0 days, 1 hours, 10 minutes, 0 seconds

To add additional VLAN groups to an MSTP instance:

1. Click Spanning Tree, MSTP.
2. Select Configure Global from the Step list.
3. Select Add Member from the Action list.
4. Select an MST instance from the MST ID list.
5. Enter the VLAN group to add to the instance in the VLAN ID field. Note that the specified member does not have to be a configured VLAN.
6. Click Apply

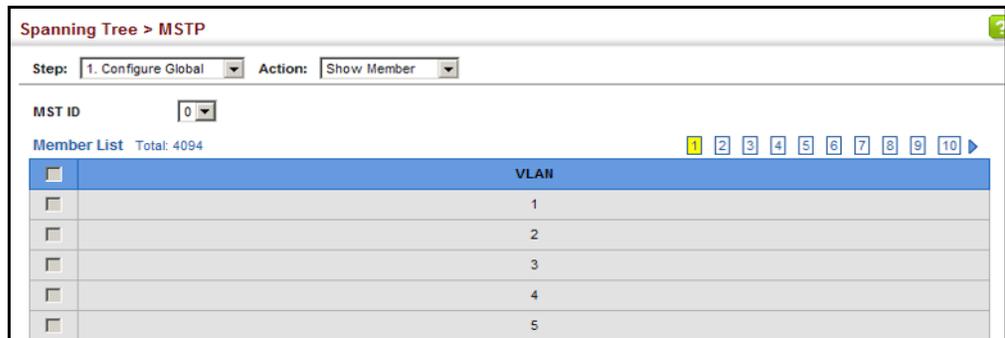
Figure 102: Adding a VLAN to an MST Instance



To show the VLAN members of an MSTP instance:

1. Click Spanning Tree, MSTP.
2. Select Configure Global from the Step list.
3. Select Show Member from the Action list.

Figure 103: Displaying Members of an MST Instance



Configuring Interface Settings for MSTP

Use the Spanning Tree > MSTP (Configure Interface - Configure) page to configure the STA interface settings for an MST instance.

Parameters

These parameters are displayed:

- ◆ **MST ID** – Instance identifier to configure. (Default: 0)
- ◆ **Interface** – Displays a list of ports or trunks.

- ◆ **STA Status** – Displays the current state of this interface within the Spanning Tree. (See “[Displaying Interface Settings for STA](#)” on page 186 for additional information.)
 - **Discarding** – Port receives STA configuration messages, but does not forward packets.
 - **Learning** – Port has transmitted configuration messages for an interval set by the Forward Delay parameter without receiving contradictory information. Port address table is cleared, and the port begins learning addresses.
 - **Forwarding** – Port forwards packets, and continues learning addresses.
- ◆ **Priority** – Defines the priority used for this port in the Spanning Tree Protocol. If the path cost for all ports on a switch are the same, the port with the highest priority (i.e., lowest value) will be configured as an active link in the Spanning Tree. This makes a port with higher priority less likely to be blocked if the Spanning Tree Protocol is detecting network loops. Where more than one port is assigned the highest priority, the port with lowest numeric identifier will be enabled. (Default: 128; Range: 0-240, in steps of 16)
- ◆ **Admin MST Path Cost** – This parameter is used by the MSTP to determine the best path between devices. Therefore, lower values should be assigned to ports attached to faster media, and higher values assigned to ports with slower media. (Path cost takes precedence over port priority.) Note that when the Path Cost Method is set to short (page 3-63), the maximum path cost is 65,535.

By default, the system automatically detects the speed and duplex mode used on each port, and configures the path cost according to the values shown below. Path cost “0” is used to indicate auto-configuration mode. When the short path cost method is selected and the default path cost recommended by the IEEE 8021w standard exceeds 65,535, the default is set to 65,535.

The recommended range is listed in [Table 11 on page 182](#).

The default path costs are listed in [Table 12 on page 183](#).

Web Interface

To configure MSTP parameters for a port or trunk:

1. Click Spanning Tree, MSTP.
2. Select Configure Interface from the Step list.
3. Select Configure from the Action list.
4. Enter the priority and path cost for an interface
5. Click Apply.

Figure 104: Configuring MSTP Interface Settings

Spanning Tree > MSTP

Step: 2. Configure Interface Action: Configure

MST ID: 0

Interface: Port Trunk

Spanning Tree Port List Total: 54

Port	STA Status	Priority (0-240, in steps of 16)	Admin MST Path Cost (0-200000000, 0: Auto)
1	Discarding	128	0
2	Discarding	128	0
3	Discarding	128	0
4	Discarding	128	0
5	Discarding	128	0

To display MSTP parameters for a port or trunk:

1. Click Spanning Tree, MSTP.
2. Select Configure Interface from the Step list.
3. Select Show Information from the Action list.

Figure 105: Displaying MSTP Interface Settings

Spanning Tree > MSTP

Step: 2. Configure Interface Action: Show Information

MST ID: 0

Interface: Port Trunk

Spanning Tree Port List Total: 54

Port	STA Status	Forward Transitions	Designated Cost	Designated Bridge	Designated Port	Oper Path Cost	Oper Link Type	Oper Edge Port	Port Role
1	Discarding	0	250	32768.0.7072CFEA1B71	128.1	1000	Point-to-Point	Disabled	Disabled
2	Discarding	0	250	32768.0.7072CFEA1B71	128.2	1000	Point-to-Point	Disabled	Disabled
3	Discarding	0	250	32768.0.7072CFEA1B71	128.3	1000	Point-to-Point	Disabled	Disabled
4	Discarding	0	250	32768.0.7072CFEA1B71	128.4	1000	Point-to-Point	Disabled	Disabled
5	Discarding	0	250	32768.0.7072CFEA1B71	128.5	1000	Point-to-Point	Disabled	Disabled

8

Congestion Control

The switch can control traffic storms by setting a maximum threshold for broadcast traffic or multicast traffic.

Congestion Control includes following options:

- ◆ **Rate Limiting** – Sets the input and output rate limits for a port.
- ◆ **Storm Control** – Sets the traffic storm threshold for each interface.

Rate Limiting

Use the Traffic > Rate Limit page to apply rate limiting to ingress or egress ports. This function allows the network manager to control the maximum rate for traffic received or transmitted on an interface. Rate limiting is configured on interfaces at the edge of a network to limit traffic into or out of the network. Packets that exceed the acceptable amount of traffic are dropped.

Rate limiting can be applied to individual ports. When an interface is configured with this feature, the traffic rate will be monitored by the hardware to verify conformity. Non-conforming traffic is dropped, conforming traffic is forwarded without any changes.

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays the switch's ports or trunks.
- ◆ **Type** – Indicates the port type. (1000BASE SFP, 10GBASE SFP+, 40GBASE QSFP)
- ◆ **Status** – Enables or disables the rate limit. (Default: Disabled)
- ◆ **Rate** – Sets the rate limit level.
(Range: 64 - 10,000,000 kbits per second for 10G Ethernet ports;
64 - 40,000,000 kbits per second for 40G Ethernet ports)

Web Interface

To configure rate limits:

1. Click Traffic, Rate Limit.
2. Set the interface type to Port or Trunk.
3. Enable the Rate Limit Status for the required ports or trunks.
4. Set the rate limit for the individual ports,.
5. Click Apply.

Figure 106: Configuring Rate Limits

Port	Type	Input		Output	
		Status	Rate (kbits/sec) (64-40000000)	Status	Rate (kbits/sec) (64-40000000)
1	10GBASE SFP+	<input type="checkbox"/> Enabled	10000000	<input type="checkbox"/> Enabled	10000000
2	10GBASE SFP+	<input type="checkbox"/> Enabled	10000000	<input type="checkbox"/> Enabled	10000000
3	10GBASE SFP+	<input type="checkbox"/> Enabled	10000000	<input type="checkbox"/> Enabled	10000000
4	10GBASE SFP+	<input type="checkbox"/> Enabled	10000000	<input type="checkbox"/> Enabled	10000000
5	10GBASE SFP+	<input type="checkbox"/> Enabled	10000000	<input type="checkbox"/> Enabled	10000000

Storm Control

Use the Traffic > Storm Control page to configure broadcast, multicast, and unknown unicast storm control thresholds. Traffic storms may occur when a device on your network is malfunctioning, or if application programs are not well designed or properly configured. If there is too much traffic on your network, performance can be severely degraded or everything can come to complete halt.

You can protect your network from traffic storms by setting a threshold for broadcast, multicast or unknown unicast traffic. Any packets exceeding the specified threshold will then be dropped.

Command Usage

- ◆ Broadcast Storm Control is enabled by default.
- ◆ When traffic exceeds the threshold specified for broadcast and multicast or unknown unicast traffic, packets exceeding the threshold are dropped until the rate falls back down beneath the threshold.
- ◆ Using both rate limiting and storm control on the same interface may lead to unexpected results. It is therefore not advisable to use both of these features on the same interface.

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Type** – Indicates interface type. (1000BASE SFP, 10GBASE SFP+, 40GBASE QSFP)
- ◆ **Unknown Unicast** – Specifies storm control for unknown unicast traffic.
- ◆ **Multicast** – Specifies storm control for multicast traffic.
- ◆ **Broadcast** – Specifies storm control for broadcast traffic.
- ◆ **Status** – Enables or disables storm control. (Default: Enabled for broadcast storm control, disabled for multicast and unknown unicast storm control)
- ◆ **Rate** – Threshold level as a rate; i.e., packets per second. (Range: 500-59520000 pps; Default: Disabled for unknown unicast and multicast traffic, 500 pps for broadcast traffic)

Web Interface

To configure broadcast storm control:

1. Click Traffic, Storm Control.
2. Set the interface type to Port or Trunk.
3. Set the Status field to enable or disable storm control.
4. Set the required threshold beyond which the switch will start dropping packets.
5. Click Apply.

Figure 107: Configuring Storm Control

Port	Type	Unknown Unicast		Multicast		Broadcast	
		Status	Rate (packets/sec) (500-59520000)	Status	Rate (packets/sec) (500-59520000)	Status	Rate (packets/sec) (500-59520000)
1	10GBASE SFP+	<input type="checkbox"/> Enabled	500	<input type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500
2	10GBASE SFP+	<input type="checkbox"/> Enabled	500	<input type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500
3	10GBASE SFP+	<input type="checkbox"/> Enabled	500	<input type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500
4	10GBASE SFP+	<input type="checkbox"/> Enabled	500	<input type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500
5	10GBASE SFP+	<input type="checkbox"/> Enabled	500	<input type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500

Class of Service

Class of Service (CoS) allows you to specify which data packets have greater precedence when traffic is buffered in the switch due to congestion. This switch supports CoS with eight priority queues for each port. Data packets in a port's high-priority queue will be transmitted before those in the lower-priority queues. You can set the default priority for each interface, and configure the mapping of frame priority tags to the switch's priority queues.

This chapter describes the following basic topics:

- ◆ [Layer 2 Queue Settings](#) – Configures each queue, including the default priority, queue mode, queue weight, and mapping of packets to queues based on CoS tags.
- ◆ [Layer 3/4 Priority Settings](#) – Selects the method by which inbound packets are processed (DSCP or CoS), and sets the per-hop behavior and drop precedence for internal processing.

Layer 2 Queue Settings

This section describes how to configure the default priority for untagged frames, set the queue mode, set the weights assigned to each queue, and map class of service tags to queues.

Setting the Default Priority for Interfaces

Use the Traffic > Priority > Default Priority page to specify the default port priority for each interface on the switch. All untagged packets entering the switch are tagged with the specified default port priority, and then sorted into the appropriate priority queue at the output port.

Command Usage

- ◆ This switch provides eight priority queues for each port. It uses Weighted Round Robin to prevent head-of-queue blockage, but can be configured to process each queue in strict order, or use a combination of strict and weighted queueing.
- ◆ The default priority applies for an untagged frame received on a port set to accept all frame types (i.e, receives both untagged and tagged frames). This priority does not apply to IEEE 802.1Q VLAN tagged frames. If the incoming frame is an IEEE 802.1Q VLAN tagged frame, the IEEE 802.1p User Priority bits will be used.

- ◆ If the output port is an untagged member of the associated VLAN, these frames are stripped of all VLAN tags prior to transmission.

Parameters

These parameters are displayed:

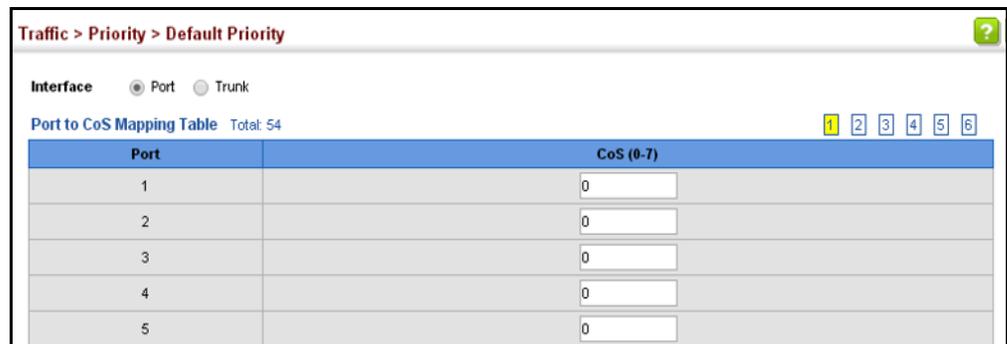
- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **CoS** – The priority that is assigned to untagged frames received on the specified interface. (Range: 0-7; Default: 0)

Web Interface

To configure the queue mode:

1. Click Traffic, Priority, Default Priority.
2. Select the interface type to display (Port or Trunk).
3. Modify the default priority for any interface.
4. Click Apply.

Figure 108: Setting the Default Port Priority



Traffic > Priority > Default Priority

Interface Port Trunk

Port to CoS Mapping Table Total: 54

Port	CoS (0-7)
1	0
2	0
3	0
4	0
5	0

Selecting the Queue Mode

Use the Traffic > Priority > Queue page to set the queue mode for the egress queues on any interface. The switch can be set to service the queues based on a strict rule that requires all traffic in a higher priority queue to be processed before the lower priority queues are serviced, or Weighted Round-Robin (WRR) queuing which specifies a scheduling weight for each queue. It can also be configured to use a combination of strict and weighted queuing.

Command Usage

- ◆ Strict priority requires all traffic in a higher priority queue to be processed before lower priority queues are serviced.
- ◆ WRR queuing specifies a relative weight for each queue. WRR uses a predefined relative weight for each queue that determines the percentage of service time

the switch services each queue before moving on to the next queue. This prevents the head-of-line blocking that can occur with strict priority queuing.

- ◆ If Strict and WRR mode is selected, a combination of strict and weighted service is used as specified for each queue. The queues assigned to use strict priority should be specified using the Strict Mode field parameter.
- ◆ A weight can be assigned to each of the weighted queues (and thereby to the corresponding traffic priorities). This weight sets the frequency at which each queue is polled for service, and subsequently affects the response time for software applications assigned a specific priority value.

Service time is shared at the egress ports by defining scheduling weights for WRR, or the queuing modes that use a combination of strict and weighted queuing. Service time is allocated to each queue by calculating a precise number of bytes per second that will be serviced on each round.

Parameters

These parameters are displayed:

- ◆ **Port** – Port or trunk identifier.
- ◆ **Queue Mode**
 - **Strict** – Services the egress queues in sequential order, transmitting all traffic in the higher priority queues before servicing lower priority queues. This ensures that the highest priority packets are always serviced first, ahead of all other traffic.
 - **WRR** – Weighted Round-Robin shares bandwidth at the egress ports by using scheduling weights, and servicing each queue in a round-robin fashion. (This is the default setting.)
 - **Strict and WRR** – Uses strict or weighted service as specified for each queue.
- ◆ **Queue ID** – The ID of the priority queue. (Range: 0-7)
- ◆ **Strict Mode** – If “Strict and WRR” mode is selected, then a combination of strict and weighted service is used as specified for each queue. Use this parameter to specify the queues assigned to use strict priority when using the strict-weighted queuing mode. (Default: Disabled)
- ◆ **Weight** – Sets a weight for each queue which is used by the WRR scheduler. (Range: 1-15; Default: Weights 1, 2, 4, 6, 8, 10, 12 and 14 are assigned to queues 0 - 7 respectively)

Web Interface

To configure the queue mode:

1. Click Traffic, Priority, Queue.
2. Select a port or trunk.
3. Set the queue mode.
4. If the weighted queue mode is selected, the queue weight can be modified if required.
5. If the queue mode that uses a combination of strict and weighted queuing is selected, the queues which are serviced first must be specified by enabling strict mode parameter in the table.
6. Click Apply.

Figure 109: Setting the Queue Mode (Strict)

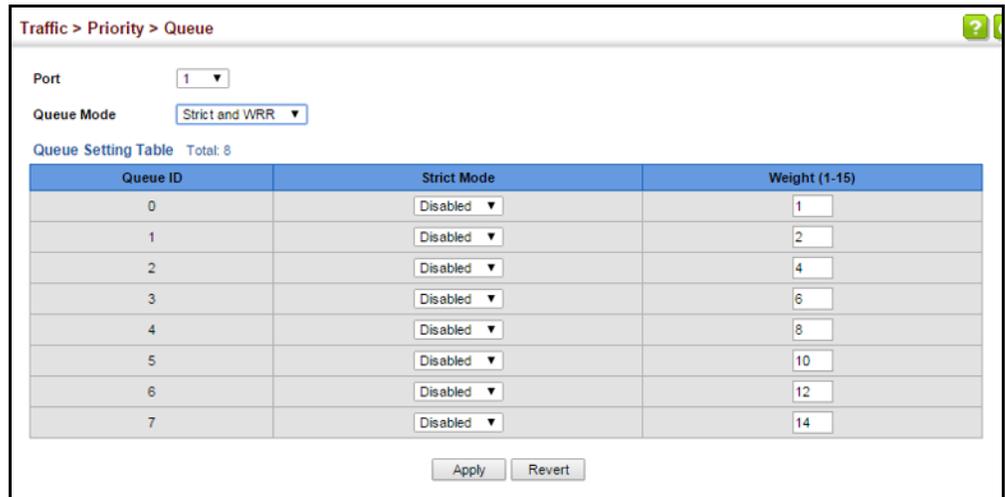
The screenshot shows the 'Traffic > Priority > Queue' configuration page. The 'Port' dropdown is set to '1'. The 'Queue Mode' dropdown is set to 'Strict'. There are 'Apply' and 'Revert' buttons at the bottom right.

Figure 110: Setting the Queue Mode (WRR)

The screenshot shows the 'Traffic > Priority > Queue' configuration page. The 'Port' dropdown is set to '1'. The 'Queue Mode' dropdown is set to 'WRR'. Below the dropdowns is a 'Queue Setting Table' with a 'Total: 8' label. The table has two columns: 'Queue ID' and 'Weight (1-15)'. The weights are set to 1, 2, 4, 6, 8, 10, 12, and 14 for queue IDs 0 through 7 respectively. There are 'Apply' and 'Revert' buttons at the bottom right.

Queue ID	Weight (1-15)
0	1
1	2
2	4
3	6
4	8
5	10
6	12
7	14

Figure 111: Setting the Queue Mode (Strict and WRR)



Mapping CoS Values to Egress Queues

Use the Traffic > Priority > PHB to Queue page to specify the hardware output queues to use based on the internal per-hop behavior value. (For more information on exact manner in which the ingress priority tags are mapped to egress queues for internal processing, see [“Mapping CoS Priorities to Internal DSCP Values” on page 210](#)).

The switch processes Class of Service (CoS) priority tagged traffic by using eight priority queues for each port, with service schedules based on strict priority, Weighted Round-Robin (WRR), or a combination of strict and weighted queuing. Up to eight separate traffic priorities are defined in IEEE 802.1p. Default priority levels are assigned according to recommendations in the IEEE 802.1p standard as shown in [Table 13](#). The following table indicates the default mapping of internal per-hop behavior to the hardware queues. The actual mapping may differ if the CoS priorities to internal DSCP values have been modified ([page 210](#)).

Table 13: IEEE 802.1p Egress Queue Priority Mapping

Priority	0	1	2	3	4	5	6	7
Queue	2	0	1	3	4	5	6	7

The priority levels recommended in the IEEE 802.1p standard for various network applications are shown in Table 14. However, priority levels can be mapped to the switch's output queues in any way that benefits application traffic for the network.

Table 14: CoS Priority Levels

Priority Level	Traffic Type
1	Background
2	(Spare)
0 (default)	Best Effort
3	Excellent Effort
4	Controlled Load
5	Video, less than 100 milliseconds latency and jitter
6	Voice, less than 10 milliseconds latency and jitter
7	Network Control

Command Usage

- ◆ Egress packets are placed into the hardware queues according to the mapping defined by this command.
- ◆ The default internal PHB to output queue mapping is shown below.

Table 15: Mapping Internal Per-hop Behavior to Hardware Queues

Per-hop Behavior	0	1	2	3	4	5	6	7
Hardware Queues	2	0	1	3	4	5	6	7

- ◆ The specified mapping applies to all interfaces.

Parameters

These parameters are displayed:

- ◆ **PHB** – Per-hop behavior, or the priority used for this router hop. (Range: 0-7, where 7 is the highest priority)
- ◆ **Queue** – Output queue buffer. (Range: 0-7, where 7 is the highest CoS priority queue)

Web Interface

To map internal PHB to hardware queues:

1. Click Traffic, Priority, PHB to Queue.
2. Select Configure from the Action list.

3. Map an internal PHB to a hardware queue. Depending on how an ingress packet is processed internally based on its CoS value, and the assigned output queue, the mapping done on this page can effectively determine the service priority for different traffic classes.
4. Click Apply.

Figure 112: Mapping CoS Values to Egress Queues

Traffic > Priority > PHB to Queue

Action:

PHB (0-7)

Queue (0-7)

To show the internal PHB to hardware queue map:

1. Click Traffic, Priority, PHB to Queue.
2. Select Show from the Action list.
3. Select an interface.

Figure 113: Showing CoS Values to Egress Queue Mapping

Traffic > Priority > PHB to Queue

Action:

PHB to Queue Mapping List Total: 8

<input type="checkbox"/>	PHB	Queue
<input type="checkbox"/>	0	2
<input type="checkbox"/>	1	0
<input type="checkbox"/>	2	1
<input type="checkbox"/>	3	3
<input type="checkbox"/>	4	4
<input type="checkbox"/>	5	5
<input type="checkbox"/>	6	6
<input type="checkbox"/>	7	7

Layer 3/4 Priority Settings

Mapping Layer 3/4 Priorities to CoS Values

The switch supports several common methods of prioritizing layer 3/4 traffic to meet application requirements. Traffic priorities can be specified in the IP header of a frame, using the priority bits in the Type of Service (ToS) octet, or the number of the TCP/UDP port. If priority bits are used, the ToS octet may contain three bits for IP Precedence or six bits for Differentiated Services Code Point (DSCP) service. When these services are enabled, the priorities are mapped to a Class of Service value by the switch, and the traffic then sent to the corresponding output queue.

Because different priority information may be contained in the traffic, this switch maps priority values to the output queues in the following manner – The precedence for priority mapping is DSCP Priority and then Default Port Priority.



Note: The default settings used for mapping priority values from ingress traffic to internal DSCP values are used to determine the hardware queues used for egress traffic, not to replace the priority values. These defaults are designed to optimize priority services for the majority of network applications. It should not be necessary to modify any of the default settings, unless a queuing problem occurs with a particular application.

Setting Priority Processing to IP Precedence/DSCP or CoS

The switch allows a choice between using IP Precedence, DSCP or CoS priority processing methods. Use the Priority > Trust Mode page to select the required processing method.

Command Usage

- ◆ If the QoS mapping mode is set to IP Precedence, and the ingress packet type is IPv4, then priority processing will be based on the IP Precedence value in the ingress packet.
- ◆ If the QoS mapping mode is set to DSCP, and the ingress packet type is IPv4, then priority processing will be based on the DSCP value in the ingress packet.
- ◆ If the QoS mapping mode is set to either IP Precedence or DSCP, and a non-IP packet is received, the packet's CoS and CFI (Canonical Format Indicator) values are used for priority processing if the packet is tagged. For an untagged packet, the default port priority (see [page 199](#)) is used for priority processing.
- ◆ If the QoS mapping mode is set to CoS, and the ingress packet type is IPv4, then priority processing will be based on the CoS and CFI values in the ingress packet.

For an untagged packet, the default port priority (see [page 199](#)) is used for priority processing.

Parameters

These parameters are displayed:

- ◆ **Port** – Specifies a port.
- ◆ **Trust Mode**
 - **CoS** – Maps layer 3/4 priorities using Class of Service values. (This is the default setting.)
 - **DSCP** – Maps layer 3/4 priorities using Differentiated Services Code Point values.
 - **IP Precedence** – Maps layer 3/4 priorities using IP Precedence values.

Web Interface

To configure the trust mode:

1. Click Traffic, Priority, Trust Mode.
2. Select the interface type to display (Port or Trunk).
3. Set the trust mode for any port.
4. Click Apply.

Figure 114: Setting the Trust Mode

Port	Trust Mode
1	CoS
2	CoS
3	CoS
4	CoS
5	CoS

Mapping Ingress DSCP Values to Internal DSCP Values

Use the Traffic > Priority > DSCP to DSCP page to map DSCP values in incoming packets to per-hop behavior and drop precedence values for internal priority processing.

The DSCP is six bits wide, allowing coding for up to 64 different forwarding behaviors. The DSCP replaces the ToS bits, but it retains backward compatibility with the three precedence bits so that non-DSCP compliant, ToS-enabled devices, will not conflict with the DSCP mapping. Based on network policies, different kinds of traffic can be marked for different kinds of forwarding.

Command Usage

- ◆ Enter per-hop behavior and drop precedence for any of the DSCP values 0 - 63.

- ◆ This map is only used when the priority mapping mode is set to DSCP (see [page 206](#)), and the ingress packet type is IPv4. Any attempt to configure the DSCP mutation map will not be accepted by the switch, unless the trust mode has been set to DSCP.
- ◆ Two QoS domains can have different DSCP definitions, so the DSCP-to-PHB/Drop Precedence mutation map can be used to modify one set of DSCP values to match the definition of another domain. The mutation map should be applied at the receiving port (ingress mutation) at the boundary of a QoS administrative domain.

Parameters

These parameters are displayed:

- ◆ **Port** – Specifies a port.
- ◆ **DSCP** – DSCP value in ingress packets. (Range: 0-63)
- ◆ **PHB** – Per-hop behavior, or the priority used for this router hop. (Range: 0-7)
- ◆ **Drop Precedence** – Drop precedence used for controlling traffic congestion. (Range: 0 - Green, 3 - Yellow, 1 - Red)

Table 16: Default Mapping of DSCP Values to Internal PHB/Drop Values

	ingress-dscp1	0	1	2	3	4	5	6	7	8	9
ingress-dscp10											
0		0,0	0,1	0,0	0,3	0,0	0,1	0,0	0,3	1,0	1,1
1		1,0	1,3	1,0	1,1	1,0	1,3	2,0	2,1	2,0	2,3
2		2,0	2,1	2,0	2,3	3,0	3,1	3,0	3,3	3,0	3,1
3		3,0	3,3	4,0	4,1	4,0	4,3	4,0	4,1	4,0	4,3
4		5,0	5,1	5,0	5,3	5,0	5,1	6,0	5,3	6,0	6,1
5		6,0	6,3	6,0	6,1	6,0	6,3	7,0	7,1	7,0	7,3
6		7,0	7,1	7,0	7,3						

The ingress DSCP is composed of ingress-dscp10 (most significant digit in the left column) and ingress-dscp1 (least significant digit in the top row (in other words, $\text{ingress-dscp} = \text{ingress-dscp10} * 10 + \text{ingress-dscp1}$); and the corresponding internal-dscp is shown at the intersecting cell in the table.

The ingress DSCP is bitwise ANDed with the binary value 11 to determine the drop precedence. If the resulting value is 10 binary, then the drop precedence is set to 0.

Web Interface

To map DSCP values to internal PHB/drop precedence:

1. Click Traffic, Priority, DSCP to DSCP.
2. Select Configure from the Action list.
3. Select a port.
4. Set the PHB and drop precedence for any DSCP value.
5. Click Apply.

Figure 115: Configuring DSCP to DSCP Internal Mapping

To show the DSCP to internal PHB/drop precedence map:

1. Click Traffic, Priority, DSCP to DSCP.
2. Select Show from the Action list.
3. Select a port.

Figure 116: Showing DSCP to DSCP Internal Mapping

	DSCP	PHB	Drop Precedence
<input type="checkbox"/>	0	0	0: Green
<input type="checkbox"/>	1	0	1: Red
<input type="checkbox"/>	2	0	0: Green
<input type="checkbox"/>	3	0	3: Yellow
<input type="checkbox"/>	4	0	0: Green
<input type="checkbox"/>	5	0	1: Red
<input type="checkbox"/>	6	0	0: Green
<input type="checkbox"/>	7	0	3: Yellow
<input type="checkbox"/>	8	1	0: Green
<input type="checkbox"/>	9	1	1: Red

Mapping CoS Priorities to Internal DSCP Values

Use the Traffic > Priority > CoS to DSCP page to maps CoS/CFI values in incoming packets to per-hop behavior and drop precedence values for priority processing.

Command Usage

- ◆ The default mapping of CoS to PHB values is shown in [Table 17 on page 210](#).
- ◆ Enter up to eight CoS/CFI paired values, per-hop behavior and drop precedence.
- ◆ If a packet arrives with a 802.1Q header but it is not an IP packet, then the CoS/CFI-to-PHB/Drop Precedence mapping table is used to generate priority and drop precedence values for internal processing. Note that priority tags in the original packet are not modified by this command.
- ◆ The internal DSCP consists of three bits for per-hop behavior (PHB) which determines the queue to which a packet is sent; and two bits for drop precedence (namely color) which is used to control traffic congestion.
- ◆ The specified mapping applies to all interfaces.

Parameters

These parameters are displayed:

- ◆ **Port** – Specifies a port.
- ◆ **CoS** – CoS value in ingress packets. (Range: 0-7)
- ◆ **CFI** – Canonical Format Indicator. Set to this parameter to “0” to indicate that the MAC address information carried in the frame is in canonical format. (Range: 0-1)
- ◆ **PHB** – Per-hop behavior, or the priority used for this router hop. (Range: 0-7)
- ◆ **Drop Precedence** – Drop precedence used for controlling traffic congestion. (Range: 0 - Green, 3 - Yellow, 1 - Red)

Table 17: Default Mapping of CoS/CFI to Internal PHB/Drop Precedence

CoS	CFI	0	1
0		(0,0)	(0,1)
1		(1,0)	(1,1)
2		(2,0)	(2,1)
3		(3,0)	(3,1)
4		(4,0)	(4,1)
5		(5,0)	(5,1)

Table 17: Default Mapping of CoS/CFI to Internal PHB/Drop Precedence

CoS	CFI	0	1
6		(6,0)	(6,1)
7		(7,0)	(7,1)

Web Interface

To map CoS/CFI values to internal PHB/drop precedence:

1. Click Traffic, Priority, CoS to DSCP.
2. Select Configure from the Action list.
3. Select a port.
4. Set the PHB and drop precedence for any of the CoS/CFI combinations.
5. Click Apply.

Figure 117: Configuring CoS to DSCP Internal Mapping

Traffic > Priority > CoS to DSCP

Action:

Port:

CoS (0-7):

CFI (0-1):

PHB (0-7):

Drop Precedence:

To show the CoS/CFI to internal PHB/drop precedence map:

1. Click Traffic, Priority, CoS to DSCP.
2. Select Show from the Action list.
3. Select a port.

Figure 118: Showing CoS to DSCP Internal Mapping

Traffic > Priority > CoS to DSCP

Action: Show

Port: 1

CoS to DSCP Mapping List Total: 16

<input type="checkbox"/>	CoS	CFI	PHB	Drop Precedence
<input type="checkbox"/>	0	0	0	0: Green
<input type="checkbox"/>	0	1	0	0: Green
<input type="checkbox"/>	1	0	1	0: Green
<input type="checkbox"/>	1	1	1	0: Green
<input type="checkbox"/>	2	0	2	0: Green
<input type="checkbox"/>	2	1	2	0: Green
<input type="checkbox"/>	3	0	3	0: Green
<input type="checkbox"/>	3	1	3	0: Green
<input type="checkbox"/>	4	0	4	0: Green
<input type="checkbox"/>	4	1	4	0: Green

Default Revert

Mapping Internal DSCP Values to Egress CoS Values

Use the Traffic > Priority > DSCP to CoS page to map internal per-hop behavior and drop precedence value pairs to CoS values used in tagged egress packets on a Layer 2 interface.

Command Usage

- ◆ Enter any per-hop behavior and drop precedence pair within the internal priority map, and then enter the corresponding CoS/CFI pair.
- ◆ If the packet is forwarded with an 8021.Q tag, the priority value in the egress packet is modified based on the default values shown in [Table 18 on page 213](#), or on the values modified by this function.

Parameters

These parameters are displayed in the web interface:

- ◆ **Port** – Port identifier.
- ◆ **PHB** – Per-hop behavior, or the priority used for this router hop. (Range: 0-7)
- ◆ **Drop Precedence** – Drop precedence used for controlling traffic congestion. (Range: 0 - Green, 3 - Yellow, 1 - Red)
- ◆ **CoS** – Class-of-Service value. (Range: 0-7)
- ◆ **CFI** – Canonical Format Indicator. Set to this parameter to “0” to indicate that the MAC address information carried in the frame is in canonical format. (Range: 0-1)

Table 18: Mapping Internal PHB/Drop Precedence to CoS/CFI Values

Drop Precedence	0 (green)	1 (red)	3 (yellow)
Per-hop Behavior			
0	(0,0)	(0,1)	(0,1)
1	(1,0)	(1,1)	(1,1)
2	(2,0)	(2,1)	(2,1)
3	(3,0)	(3,1)	(3,1)
4	(4,0)	(4,1)	(4,1)
5	(5,0)	(5,1)	(5,1)
6	(6,0)	(6,1)	(6,1)
7	(7,0)	(7,1)	(7,1)

Web Interface

To map internal per-hop behavior and drop precedence values to CoS values in the web interface:

1. Click Traffic, Priority, DSCP to CoS.
2. Select Configure from the Action list.
3. Select an interface.
4. Select any PHB and drop precedence pair within the internal priority map, and then set the corresponding CoS/CFI pair.
5. Click Apply.

Figure 119: Configuring DSCP to CoS Egress Mapping

Traffic > Priority > DSCP to CoS

Action:

Port:

PHB (0-7):

Drop Precedence:

CoS (0-7):

CFI (0-1):

To show the DSCP to CoS egress map in the web interface:

1. Click Traffic, Priority, DSCP to CoS.
2. Select Show from the Action list.

3. Select an interface.

Figure 120: Showing DSCP to CoS Egress Mapping

PHB	Drop Precedence	CoS	CFI
0	0	0	0
0	1	0	1
0	3	0	1
1	0	1	0
1	1	1	1
1	3	1	1

Mapping IP Precedence Values to Internal DSCP Values

Use the Traffic > Priority > IP Precedence to DSCP page to map IP precedence values in incoming packets to per-hop behavior and drop precedence values for priority processing.

The Type of Service (ToS) octet in the IPv4 header includes three precedence bits defining eight different priority levels ranging from highest priority for network control packets to lowest priority for routine traffic. The default IP Precedence values map one-to-one to the Class of Service values (that is, Precedence value 0 maps to PHB value 0, and so forth). Bits 6 and 7 are used for network control, and the other bits for various application types. The ToS bits are defined in Table 19.

Table 19: Mapping IP Precedence

Priority Level	Traffic Type
7	Network Control
6	Internetwork Control
5	Critical
4	Flash Override
3	Flash
2	Immediate
1	Priority
0	Routine

Command Usage

- ◆ Enter per-hop behavior and drop precedence for any of the IP Precedence values 0 - 7.

- ◆ If the priority mapping mode is set the IP Precedence and the ingress packet type is IPv4, then the IP Precedence-to-PHB/Drop Precedence mapping table is used to generate priority and drop precedence values for internal processing.

Parameters

These parameters are displayed in the web interface:

- ◆ **Port** – Port identifier.
- ◆ **IP Precedence** – IP Precedence value in ingress packets. (Range: 0-7)
- ◆ **PHB** – Per-hop behavior, or the priority used for this router hop. (Range: 0-7)
- ◆ **Drop Precedence** – Drop precedence used for controlling traffic congestion. (Range: 0 - Green, 3 - Yellow, 1 - Red)

Table 20: Default Mapping of IP Precedence to Internal PHB/Drop Values

IP Precedence Value	0	1	2	3	4	5	6	7
Per-hop Behavior	0	1	2	3	4	5	6	7
Drop Precedence	0	0	0	0	0	0	0	0

Web Interface

To map IP Precedence to internal PHB/drop precedence in the web interface:

1. Click Traffic, Priority, IP Precedence to DSCP.
2. Select Configure from the Action list.
3. Select a port.
4. Set the PHB and drop precedence for any of the IP Precedence values.
5. Click Apply.

Figure 121: Configuring IP Precedence to DSCP Internal Mapping

Traffic > Priority > IP Precedence to DSCP

Action:

Port:

IP Precedence (0-7):

PHB (0-7):

Drop Precedence:

To show the IP Precedence to internal PHB/drop precedence map in the web interface:

1. Click Traffic, Priority, IP Precedence to DSCP.
2. Select Show from the Action list.
3. Select a port.

Figure 122: Showing the IP Precedence to DSCP Internal Map

IP Precedence	PHB	Drop Precedence
0	0	0: Green
1	1	0: Green
2	2	0: Green
3	3	0: Green
4	4	0: Green
5	5	0: Green
6	6	0: Green
7	7	0: Green

Mapping IP Port Priority to Internal DSCP Values

Use the Traffic > Priority > IP Port to DSCP page to map network applications designated by a TCP/UDP destination port number in the frame header to per-hop behavior and drop precedence values for internal priority processing.

Command Usage

- ◆ This mapping table is only used if the protocol type of the arriving packet is TCP or UDP. Some of the more common TCP service ports include: HTTP: 80, FTP: 21, Telnet: 23 and POP3: 110.
- ◆ No default mapping is defined for ingress TCP/UDP port types.

Parameters

These parameters are displayed in the web interface:

- ◆ **Port** – Port identifier.
- ◆ **IP Protocol**
 - **TCP** – Transport Control Protocol
 - **UDP** – User Datagram Protocol
- ◆ **Destination Port Number** – 16-bit TCP/UDP destination port number. (Range: 0-65535)

- ◆ **PHB** – Per-hop behavior, or the priority used for this router hop. (Range: 0-7)
- ◆ **Drop Precedence** – Drop precedence used for controlling traffic congestion. (Range: 0 - Green, 3 - Yellow, 1 - Red)

Web Interface

To map TCP/UDP port number to per-hop behavior and drop precedence in the web interface:

1. Click Traffic, Priority, IP Port to DSCP.
2. Select Configure from the Action list.
3. Select a port.
4. Set the PHB and drop precedence for any TCP or UDP port.
5. Click Apply.

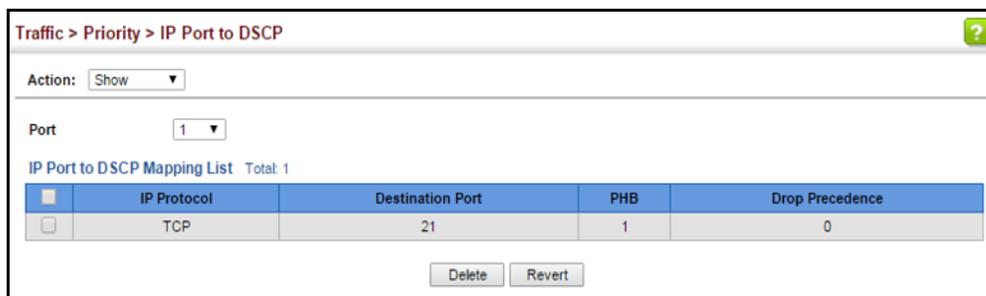
Figure 123: Configuring IP Port Number to DSCP Internal Mapping

The screenshot shows a web interface for configuring IP Port Number to DSCP Internal Mapping. The breadcrumb navigation is "Traffic > Priority > IP Port to DSCP". The "Action" dropdown is set to "Configure". The "Port" dropdown is set to "1". The "IP Protocol" dropdown is set to "TCP". The "Destination Port (0-65535)" text input field contains "21". The "PHB (0-7)" text input field contains "1". The "Drop Precedence" dropdown is set to "0: Green". There are "Apply" and "Revert" buttons at the bottom right.

To show the TCP/UDP port number to per-hop behavior and drop precedence map in the web interface:

1. Click Traffic, Priority, IP Port to DSCP.
2. Select Show from the Action list.
3. Select a port.

Figure 124: Showing IP Port Number to DSCP Internal Mapping



Quality of Service

This chapter describes the following tasks required to apply QoS policies:

Class Map – Creates a map which identifies a specific class of traffic.

Policy Map – Sets the boundary parameters used for monitoring inbound traffic, and the action to take for conforming and non-conforming traffic.

Binding to a Port – Applies a policy map to an ingress port.

Overview

The commands described in this section are used to configure Quality of Service (QoS) classification criteria and service policies. Differentiated Services (DiffServ) provides policy-based management mechanisms used for prioritizing network resources to meet the requirements of specific traffic types on a per hop basis. Each packet is classified upon entry into the network based on access lists, IP Precedence, DSCP values, VLAN lists, or CoS values. Using access lists allows you select traffic based on Layer 2, Layer 3, or Layer 4 information contained in each packet. Based on configured network policies, different kinds of traffic can be marked for different kinds of forwarding.

All switches or routers that access the Internet rely on class information to provide the same forwarding treatment to packets in the same class. Class information can be assigned by end hosts, or switches or routers along the path. Priority can then be assigned based on a general policy, or a detailed examination of the packet. However, note that detailed examination of packets should take place close to the network edge so that core switches and routers are not overloaded.

Switches and routers along the path can use class information to prioritize the resources allocated to different traffic classes. The manner in which an individual device handles traffic in the DiffServ architecture is called per-hop behavior. All devices along a path should be configured in a consistent manner to construct a consistent end-to-end QoS solution.



Note: You can configure up to 16 rules per class map. You can also include multiple classes in a policy map.

Note: You should create a class map before creating a policy map. Otherwise, you will not be able to select a class map from the policy rule settings screen (see [page 224](#)).

Command Usage

To create a service policy for a specific category or ingress traffic, follow these steps:

1. Use the Configure Class (Add) page to designate a class name for a specific category of traffic.
2. Use the Configure Class (Add Rule) page to edit the rules for each class which specify a type of traffic based on an access list, a DSCP or IP Precedence value, a VLAN or a CoS value.
3. Use the Configure Policy (Add) page to designate a policy name for a specific manner in which ingress traffic will be handled.
4. Use the Configure Policy (Add Rule) page to add one or more classes to the policy map. Assign policy rules to each class by “setting” the QoS value (CoS or PHB) to be assigned to the matching traffic class. The policy rule can also be configured to monitor the maximum throughput and burst rate. Then specify the action to take for conforming traffic, or the action to take for a policy violation.
5. Use the Configure Interface page to assign a policy map to a specific interface.



Note: Up to 32 classes can be created, but a policy map can contain only one class.

Configuring a Class Map

A class map is used for matching packets to a specified class. Use the Traffic > DiffServ (Configure Class) page to configure a class map.

Command Usage

- ◆ The class map is used with a policy map ([page 224](#)) to create a service policy ([page 233](#)) for a specific interface that defines packet classification, service tagging, and bandwidth policing. Note that one or more class maps can be assigned to a policy map.
- ◆ Up to 32 class maps can be configured.

Parameters

These parameters are displayed:

Add

- ◆ **Class Name** – Name of the class map. (Range: 1-32 characters)

- ◆ **Type** – Only one match command is permitted per class map, so the match-any field refers to the criteria specified by the lone match command.
- ◆ **Description** – A brief description of a class map. (Range: 1-64 characters)

Add Rule

- ◆ **Class Name** – Name of the class map.
- ◆ **Type** – Only one match command is permitted per class map, so the match-any field refers to the criteria specified by the lone match command.
- ◆ **ACL** – Name of an access control list. Any type of ACL can be specified, including standard or extended IPv4/IPv6 ACLs and MAC ACLs.
- ◆ **IP DSCP** – A DSCP value. (Range: 0-63)
- ◆ **IP Precedence** – An IP Precedence value. (Range: 0-7)
- ◆ **VLAN ID** – A VLAN. (Range:1-4094)
- ◆ **CoS** – A CoS value. (Range: 0-7)

Web Interface

To configure a class map:

1. Click Traffic, DiffServ.
2. Select Configure Class from the Step list.
3. Select Add from the Action list.
4. Enter a class name.
5. Enter a description.
6. Click Add.

Figure 125: Configuring a Class Map

Traffic > DiffServ

Step: 1. Configure Class Action: Add

Class Name: rd-class

Type: Match Any

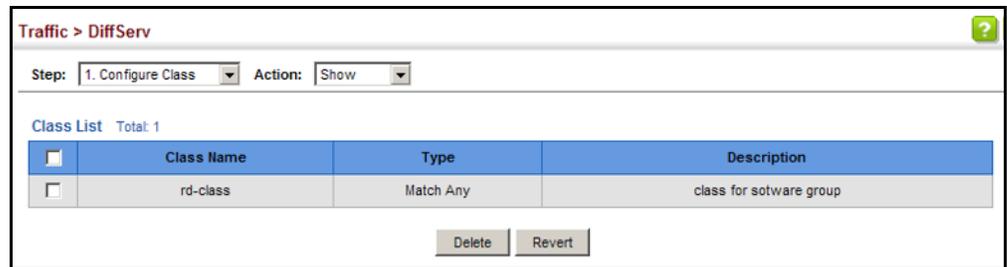
Description: class for software group

Apply Revert

To show the configured class maps:

1. Click Traffic, DiffServ.
2. Select Configure Class from the Step list.
3. Select Show from the Action list.

Figure 126: Showing Class Maps



To edit the rules for a class map:

1. Click Traffic, DiffServ.
2. Select Configure Class from the Step list.
3. Select Add Rule from the Action list.
4. Select the name of a class map.
5. Specify type of traffic for this class based on an access list, a DSCP or IP Precedence value, a VLAN, or a CoS value. You can specify up to 16 items to match when assigning ingress traffic to a class map.
6. Click Apply.

Figure 127: Adding Rules to a Class Map

The screenshot shows the configuration interface for DiffServ. The breadcrumb is "Traffic > DiffServ". The "Step" dropdown is set to "1. Configure Class" and the "Action" dropdown is set to "Add Rule". The "Class Name" is "rd" and the "Type" is "Match Any". Under the "Rule:" section, the "IP DSCP (0-63)" radio button is selected, and the value "3" is entered in the adjacent text box. Other rule options include ACL, IP Precedence (0-7), VLAN ID (1-4094), and CoS (0-7). "Apply" and "Revert" buttons are at the bottom right.

To show the rules for a class map:

1. Click Traffic, DiffServ.
2. Select Configure Class from the Step list.
3. Select Show Rule from the Action list.

Figure 128: Showing the Rules for a Class Map

The screenshot shows the configuration interface for DiffServ. The breadcrumb is "Traffic > DiffServ". The "Step" dropdown is set to "1. Configure Class" and the "Action" dropdown is set to "Show Rule". The "Class Name" is "rd-class" and the "Type" is "Match Any". Below this is a "Rule List" table with a "Total: 2" count. The table has two rows, each with a checkbox and a "Rule" column. The first row shows "IP DSCP 3" and the second row shows "IP Precedence 3". "Delete" and "Revert" buttons are at the bottom right.

Rule List Total: 2	
	Rule
<input type="checkbox"/>	IP DSCP 3
<input type="checkbox"/>	IP Precedence 3

Creating QoS Policies

Use the Traffic > DiffServ (Configure Policy) page to create a policy map that can be attached to multiple interfaces. A policy map is used to group one or more class map statements (page 220), modify service tagging, and enforce bandwidth policing. A policy map can then be bound by a service policy to one or more interfaces (page 233).

Configuring QoS policies requires several steps. A class map must first be configured which indicates how to match the inbound packets according to an access list, a DSCP or IP Precedence value, or a member of specific VLAN. A policy map is then configured which indicates the boundary parameters used for monitoring inbound traffic, and the action to take for conforming and non-conforming traffic. A policy map may contain one or more classes based on previously defined class maps.

The class of service or per-hop behavior (i.e., the priority used for internal queue processing) can be assigned to matching packets. In addition, the flow rate of inbound traffic can be monitored and the response to conforming and non-conforming traffic based by one of three distinct policing methods as described below.

Police Flow Meter – Defines the committed information rate (maximum throughput), committed burst size (burst rate), and the action to take for conforming and non-conforming traffic.

Policing is based on a token bucket, where bucket depth (that is, the maximum burst before the bucket overflows) is specified by the “burst” field (BC), and the average rate tokens are removed from the bucket is specified by the “rate” option (CIR). Action may be taken for traffic conforming to the maximum throughput, or exceeding the maximum throughput.

srTCM Police Meter – Defines an enforcer for classified traffic based on a single rate three color meter scheme defined in RFC 2697. This metering policy monitors a traffic stream and processes its packets according to the committed information rate (CIR, or maximum throughput), committed burst size (BC, or burst rate), and excess burst size (BE). Action may taken for traffic conforming to the maximum throughput, exceeding the maximum throughput, or exceeding the excess burst size.

- ◆ The PHB label is composed of five bits, three bits for per-hop behavior, and two bits for the color scheme used to control queue congestion. In addition to the actions defined by this command to transmit, remark the DSCP service value, or drop a packet, the switch will also mark the two color bits used to set the drop precedence of a packet. A packet is marked green if it doesn't exceed the committed information rate and committed burst size, yellow if it does exceed the committed information rate and committed burst size, but not the excess burst size, and red otherwise.

- ◆ The meter operates in one of two modes. In the color-blind mode, the meter assumes that the packet stream is uncolored. In color-aware mode the meter assumes that some preceding entity has pre-colored the incoming packet stream so that each packet is either green, yellow, or red. The marker (re)colors an IP packet according to the results of the meter. The color is coded in the DS field [RFC 2474] of the packet.
- ◆ The behavior of the meter is specified in terms of its mode and two token buckets, C and E, which both share the common rate CIR. The maximum size of the token bucket C is BC and the maximum size of the token bucket E is BE.

The token buckets C and E are initially full, that is, the token count $T_c(0) = BC$ and the token count $T_e(0) = BE$. Thereafter, the token counts T_c and T_e are updated CIR times per second as follows:

- If T_c is less than BC, T_c is incremented by one, else
- if T_e is less than BE, T_e is incremented by one, else
- neither T_c nor T_e is incremented.

When a packet of size B bytes arrives at time t, the following happens if srTCM is configured to operate in Color-Blind mode:

- If $T_c(t) - B \geq 0$, the packet is green and T_c is decremented by B down to the minimum value of 0, else
- if $T_e(t) - B \geq 0$, the packets is yellow and T_e is decremented by B down to the minimum value of 0,
- else the packet is red and neither T_c nor T_e is decremented.

When a packet of size B bytes arrives at time t, the following happens if srTCM is configured to operate in Color-Aware mode:

- If the packet has been precolored as green and $T_c(t) - B \geq 0$, the packet is green and T_c is decremented by B down to the minimum value of 0, else
- If the packet has been precolored as yellow or green and if $T_e(t) - B \geq 0$, the packets is yellow and T_e is decremented by B down to the minimum value of 0, else
- the packet is red and neither T_c nor T_e is decremented.

The metering policy guarantees a deterministic behavior where the volume of green packets is never smaller than what has been determined by the CIR and BC, that is, tokens of a given color are always spent on packets of that color. Refer to RFC 2697 for more information on other aspects of srTCM.

trTCM Police Meter – Defines an enforcer for classified traffic based on a two rate three color meter scheme defined in RFC 2698. This metering policy monitors a traffic stream and processes its packets according to the committed information rate (CIR, or maximum throughput), peak information rate (PIR), and their associated burst sizes – committed burst size (BC, or burst rate), and peak burst size

(BP). Action may taken for traffic conforming to the maximum throughput, exceeding the maximum throughput, or exceeding the peak burst size.

- ◆ The PHB label is composed of five bits, three bits for per-hop behavior, and two bits for the color scheme used to control queue congestion. In addition to the actions defined by this command to transmit, remark the DSCP service value, or drop a packet, the switch will also mark the two color bits used to set the drop precedence of a packet. A packet is marked red if it exceeds the PIR. Otherwise it is marked either yellow or green depending on whether it exceeds or doesn't exceed the CIR.

The trTCM is useful for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate.

- ◆ The meter operates in one of two modes. In the color-blind mode, the meter assumes that the packet stream is uncolored. In color-aware mode the meter assumes that some preceding entity has pre-colored the incoming packet stream so that each packet is either green, yellow, or red. The marker (re)colors an IP packet according to the results of the meter. The color is coded in the DS field [RFC 2474] of the packet.
- ◆ The behavior of the meter is specified in terms of its mode and two token buckets, P and C, which are based on the rates PIR and CIR, respectively. The maximum size of the token bucket P is BP and the maximum size of the token bucket C is BC.

The token buckets P and C are initially (at time 0) full, that is, the token count $Tp(0) = BP$ and the token count $Tc(0) = BC$. Thereafter, the token count Tp is incremented by one PIR times per second up to BP and the token count Tc is incremented by one CIR times per second up to BC.

When a packet of size B bytes arrives at time t, the following happens if trTCM is configured to operate in Color-Blind mode:

- If $Tp(t)-B < 0$, the packet is red, else
- if $Tc(t)-B < 0$, the packet is yellow and Tp is decremented by B, else
- the packet is green and both Tp and Tc are decremented by B.

When a packet of size B bytes arrives at time t, the following happens if trTCM is configured to operate in Color-Aware mode:

- If the packet has been precolored as red or if $Tp(t)-B < 0$, the packet is red, else
 - if the packet has been precolored as yellow or if $Tc(t)-B < 0$, the packet is yellow and Tp is decremented by B, else
 - the packet is green and both Tp and Tc are decremented by B.
- ◆ The trTCM can be used to mark a IP packet stream in a service, where different, decreasing levels of assurances (either absolute or relative) are given to packets

which are green, yellow, or red. Refer to RFC 2698 for more information on other aspects of trTCM.

Command Usage

- ◆ A policy map can contain 16 class statements that can be applied to the same interface (page 233). Up to 32 policy maps can be configured for ingress ports.
- ◆ After using the policy map to define packet classification, service tagging, and bandwidth policing, it must be assigned to a specific interface by a service policy (page 233) to take effect.

Parameters

These parameters are displayed:

Add

- ◆ **Policy Name** – Name of policy map. (Range: 1-32 characters)
- ◆ **Description** – A brief description of a policy map. (Range: 1-64 characters)

Add Rule

- ◆ **Policy Name** – Name of policy map.
- ◆ **Class Name** – Name of a class map that defines a traffic classification upon which a policy can act. A policy map can contain only one class map.
- ◆ **Action** – This attribute is used to set an internal QoS value in hardware for matching packets. The PHB label is composed of five bits, three bits for per-hop behavior, and two bits for the color scheme used to control queue congestion with the srTCM and trTCM metering functions.
 - **Set CoS** – Configures the service provided to ingress traffic by setting an internal CoS value for a matching packet (as specified in rule settings for a class map). (Range: 0-7)
See Table 17, “Default Mapping of CoS/CFI to Internal PHB/Drop Precedence,” on page 210).
 - **Set PHB** – Configures the service provided to ingress traffic by setting the internal per-hop behavior for a matching packet (as specified in rule settings for a class map). (Range: 0-7)
See Table 17, “Default Mapping of CoS/CFI to Internal PHB/Drop Precedence,” on page 210).
- ◆ **Meter** – Check this to define the maximum throughput, burst rate, and the action that results from a policy violation.

- ◆ **Meter Mode** – Selects one of the following policing methods.
 - **Flow** (Police Flow) – Defines the committed information rate (CIR, or maximum throughput), committed burst size (BC, or burst rate), and the action to take for conforming and non-conforming traffic. Policing is based on a token bucket, where bucket depth (that is, the maximum burst before the bucket overflows) is specified by the “burst” field, and the average rate tokens are removed from the bucket is by specified by the “rate” option.
 - **Committed Information Rate** (CIR) – Committed information rate in kilobits per second. (Range: 0-40000000 kbps or maximum port speed, whichever is lower)
The rate cannot exceed the configured interface speed.
 - **Committed Burst Size** (BC) – Committed burst in bytes. (Range: 1000-128000000 bytes)
 - **Conform** – Specifies that traffic conforming to the committed information rate (CIR) and committed burst size (BC) will be transmitted without any change to the DSCP service level.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Transmit** – Transmits in-conformance traffic without any change to the DSCP service level.
 - **Violate** – Specifies whether the traffic that exceeds the committed information rate (CIR) or burst size (BC) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Drop** – Drops out of conformance traffic.
 - **srTCM** (Police Meter) – Defines the committed information rate (CIR, or maximum throughput), committed burst size (BC, or burst rate) and excess burst size (BE), and the action to take for traffic conforming to the maximum throughput, exceeding the maximum throughput but within the excess burst size, or exceeding the excess burst size. In addition to the actions defined by this command to transmit, remark the DSCP service value, or drop a packet, the switch will also mark the two color bits used to set the drop precedence of a packet.

The color modes include “Color-Blind” which assumes that the packet stream is uncolored, and “Color-Aware” which assumes that the incoming packets are pre-colored. The functional differences between these modes is described at the beginning of this section under “srTCM Police Meter.”

- **Committed Information Rate (CIR)** – Committed information rate in kilobits per second. (Range: 0-40000000 kbps at a granularity of 64 kbps or maximum port speed, whichever is lower)
The rate cannot exceed the configured interface speed.
- **Committed Burst Size (BC)** – Committed burst size in bytes. (Range: 1000-128000000 bytes at a granularity of 4k bytes)
- **Excess Burst Size (BE)** – Burst in excess of committed burst size. (Range: 1000-128000000 bytes at a granularity of 4k bytes)
- **Conform** – Specifies that traffic conforming to the maximum committed information rate (CIR) and committed burst size (BC) will be transmitted without any change to the DSCP service level.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Transmit** – Transmits in-conformance traffic without any change to the DSCP service level.
- **Exceed** – Specifies whether traffic that exceeds the maximum rate (CIR) but is within the excess burst size (BE) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Drop** – Drops out of conformance traffic.
- **Violate** – Specifies whether the traffic that exceeds the excess burst size (BE) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Drop** – Drops out of conformance traffic.
- **trTCM (Police Meter)** – Defines the committed information rate (CIR, or maximum throughput), peak information rate (PIR), and their associated burst sizes – committed burst size (BC, or burst rate) and peak burst size (BP), and the action to take for traffic conforming to the maximum throughput, exceeding the maximum throughput but within the peak information rate, or exceeding the peak information rate. In addition to the actions defined by this command to transmit, remark the DSCP service value, or drop a packet, the switch will also mark the two color bits used to set the drop precedence of a packet.

The color modes include “Color-Blind” which assumes that the packet stream is uncolored, and “Color-Aware” which assumes that the incoming packets are pre-colored. The functional differences between these modes is described at the beginning of this section under “trTCM Police Meter.”

- **Committed Information Rate (CIR)** – Committed information rate in kilobits per second. (Range: 0-40000000 kbps at a granularity of 64 kbps or maximum port speed, whichever is lower)
The rate cannot exceed the configured interface speed.
- **Committed Burst Size (BC)** – Committed burst size in bytes. (Range: 1000-128000000 bytes at a granularity of 4k bytes)
- **Peak Information Rate (PIR)** – Rate in kilobits per second. (Range: 0-40000000 kbps at a granularity of 64 kbps or maximum port speed, whichever is lower)
The rate cannot exceed the configured interface speed.
- **Peak Burst Size (BP)** – Committed burst size in bytes. (Range: 1000-128000000 bytes at a granularity of 4k bytes)
- **Conform** – Specifies that traffic conforming to the committed maximum rate (CIR) and peak burst size (BP) will be transmitted without any change to the DSCP service level.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63).
 - **Transmit** – Transmits in-conformance traffic without any change to the DSCP service level.
- **Exceed** – Specifies whether traffic that exceeds the maximum rate (CIR) but is within the peak information rate (PIR) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63).
 - **Drop** – Drops out of conformance traffic.
- **Violate** – Specifies whether the traffic that exceeds the peak information rate (PIR) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63).
 - **Drop** – Drops out of conformance traffic.

Web Interface

To configure a policy map:

1. Click Traffic, DiffServ.
2. Select Configure Policy from the Step list.
3. Select Add from the Action list.
4. Enter a policy name.
5. Enter a description.
6. Click Add.

Figure 129: Configuring a Policy Map

Traffic > DiffServ

Step: 2. Configure Policy Action: Add

Policy Name rd-policy

Description for the software group

Apply Revert

To show the configured policy maps:

1. Click Traffic, DiffServ.
2. Select Configure Policy from the Step list.
3. Select Show from the Action list.

Figure 130: Showing Policy Maps

Traffic > DiffServ

Step: 2. Configure Policy Action: Show

Policy List Total: 1

	Policy Name	Description
<input type="checkbox"/>	rd-policy	for the software group

Delete Revert

To edit the rules for a policy map:

1. Click Traffic, DiffServ.
2. Select Configure Policy from the Step list.
3. Select Add Rule from the Action list.
4. Select the name of a policy map.
5. Set the CoS or per-hop behavior for matching packets to specify the quality of service to be assigned to the matching traffic class. Use one of the metering options to define parameters such as the maximum throughput and burst rate. Then specify the action to take for conforming traffic, the action to take for traffic in excess of the maximum rate but within the peak information rate, or the action to take for a policy violation.
6. Click Apply.

Figure 131: Adding Rules to a Policy Map

The screenshot shows the configuration interface for adding a rule to a DiffServ policy map. The breadcrumb is "Traffic > DiffServ". The "Step" is "2. Configure Policy" and the "Action" is "Add Rule". The "Policy Name" is "rd-policy".

Rule:

- Class Name:** rd-class
- Action:** Set PHB (0-7) 3
- Meter:**

Meter Mode: Flow

Committed Information Rate (0-40000000): 1000000 kbps

Committed Burst Size (1000-128000000): 4000 bytes

Excess Burst Size (1000-128000000): bytes

Peak Information Rate (0-40000000): kbps

Peak Burst Size (1000-128000000): bytes

Conform: Transmit

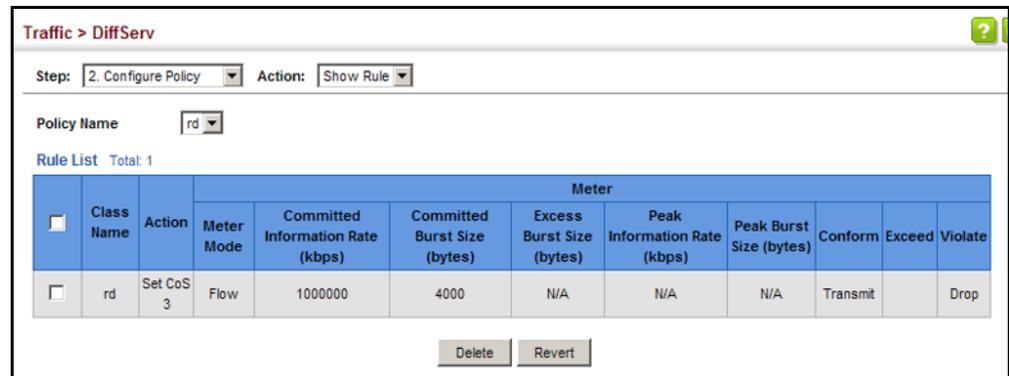
Exceed: Set IP DSCP (0-63)

Violate: Drop

To show the rules for a policy map:

1. Click Traffic, DiffServ.
2. Select Configure Policy from the Step list.
3. Select Show Rule from the Action list.

Figure 132: Showing the Rules for a Policy Map



Attaching a Policy Map to a Port

Use the Traffic > DiffServ (Configure Interface) page to bind a policy map to a port.

Command Usage

- ◆ First define a class map, define a policy map, and bind the service policy to the required interface.
- ◆ The switch does not allow a policy map to be bound to an interface for egress traffic.

Parameters

These parameters are displayed:

- ◆ **Port** – Specifies a port.
- ◆ **Ingress** – Applies the selected rule to ingress traffic.
- ◆ **Egress** – Applies the selected rule to egress traffic.

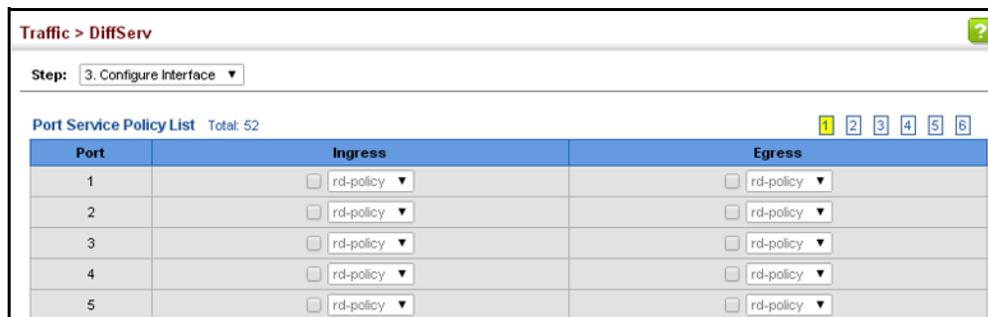
Web Interface

To bind a policy map to a port:

1. Click Traffic, DiffServ.
2. Select Configure Interface from the Step list.

3. Check the box under the Ingress or Egress field to enable a policy map for a port.
4. Select a policy map from the scroll-down box.
5. Click Apply.

Figure 133: Attaching a Policy Map to a Port



The screenshot shows a configuration page titled "Traffic > DiffServ" with a "Step: 3. Configure Interface" dropdown. Below is a "Port Service Policy List" with a total of 52 items. The table has three columns: "Port", "Ingress", and "Egress". Each row represents a port configuration with checkboxes and dropdown menus for selecting a policy map.

Port	Ingress	Egress
1	<input type="checkbox"/> rd-policy ▼	<input type="checkbox"/> rd-policy ▼
2	<input type="checkbox"/> rd-policy ▼	<input type="checkbox"/> rd-policy ▼
3	<input type="checkbox"/> rd-policy ▼	<input type="checkbox"/> rd-policy ▼
4	<input type="checkbox"/> rd-policy ▼	<input type="checkbox"/> rd-policy ▼
5	<input type="checkbox"/> rd-policy ▼	<input type="checkbox"/> rd-policy ▼

Security Measures

You can configure this switch to authenticate users logging into the system for management access using local or remote authentication methods. Port-based authentication using IEEE 802.1X can also be configured to control either management access to the uplink ports or client access to the data ports. This switch provides secure network management access using the following options:

- ◆ **AAA** – Use local or remote authentication to configure access rights, and specify authentication servers.
- ◆ **User Accounts** – Manually configure access rights on the switch for specified users.
- ◆ **Web Authentication** – Allows stations to authenticate and access the network in situations where 802.1X or Network Access authentication methods are infeasible or impractical.
- ◆ **Network Access** - Configure MAC authentication, intrusion response, dynamic VLAN assignment, and dynamic QoS assignment.
- ◆ **HTTPS** – Provide a secure web connection.
- ◆ **SSH** – Provide a secure shell (for secure Telnet access).
- ◆ **ACL** – Access Control Lists provide packet filtering for IP frames (based on address, protocol, Layer 4 protocol port number or TCP control code).
- ◆ **ARP Inspection** – Security feature that validates the MAC Address bindings for Address Resolution Protocol packets. Provides protection against ARP traffic with invalid MAC to IP Address bindings, which forms the basis for certain “man-in-the-middle” attacks.
- ◆ **IP Filter** – Filters management access to the web, SNMP or Telnet interface.
- ◆ **Port Security** – Configure secure addresses for individual ports.
- ◆ **Port Authentication** – Use IEEE 802.1X port authentication to control access to specific ports.
- ◆ **IPv4 Source Guard** – Filters IPv4 traffic on insecure ports for which the source address cannot be identified via DHCPv4 snooping nor static source bindings.
- ◆ **IPv6 Source Guard** – Filters IPv6 traffic on insecure ports for which the source address cannot be identified via ND snooping, DHCPv6 snooping, nor static source bindings.

- ◆ **DHCP Snooping** – Filter IP traffic on insecure ports for which the source address cannot be identified via DHCP snooping.



Note: The priority of execution for the filtering commands is Port Security, Port Authentication, Network Access, Web Authentication, Access Control Lists, IP Source Guard, and then DHCP Snooping.

AAA Authentication, Authorization and Accounting

The authentication, authorization, and accounting (AAA) feature provides the main framework for configuring access control on the switch. The three security functions can be summarized as follows:

- ◆ **Authentication** — Identifies users that request access to the network.
- ◆ **Authorization** — Determines if users can access specific services.
- ◆ **Accounting** — Provides reports, auditing, and billing for services that users have accessed on the network.

The AAA functions require the use of configured RADIUS or TACACS+ servers in the network. The security servers can be defined as sequential groups that are applied as a method for controlling user access to specified services. For example, when the switch attempts to authenticate a user, a request is sent to the first server in the defined group, if there is no response the second server will be tried, and so on. If at any point a pass or fail is returned, the process stops.

The switch supports the following AAA features:

- ◆ Authentication of users that access management interfaces on the switch through the console and Telnet.

To configure AAA on the switch, you need to follow this general process:

1. Configure RADIUS and TACACS+ server access parameters. See [“Configuring Local/Remote Logon Authentication” on page 237](#).
2. Define RADIUS and TACACS+ server groups to support the accounting and authorization of services.
3. Define a method name for each service to which you want to apply accounting or authorization and specify the RADIUS or TACACS+ server groups to use.
4. Apply the method names to port or line interfaces.



Note: This guide assumes that RADIUS and TACACS+ servers have already been configured to support AAA. The configuration of RADIUS and TACACS+ server software is beyond the scope of this guide, refer to the documentation provided with the RADIUS or TACACS+ server software.

Configuring Local/ Remote Logon Authentication

Use the Security > AAA > System Authentication page to specify local or remote authentication. Local authentication restricts management access based on user names and passwords manually configured on the switch. Remote authentication uses a remote access authentication server based on RADIUS or TACACS+ protocols to verify management access.

Command Usage

- ◆ By default, management access is always checked against the authentication database stored on the local switch. If a remote authentication server is used, you must specify the authentication sequence. Then specify the corresponding parameters for the remote authentication protocol using the Security > AAA > Server page. Local and remote logon authentication control management access via the console port, web browser, or Telnet.
- ◆ You can specify up to three authentication methods for any user to indicate the authentication sequence. For example, if you select (1) RADIUS, (2) TACACS and (3) Local, the user name and password on the RADIUS server is verified first. If the RADIUS server is not available, then authentication is attempted using the TACACS+ server, and finally the local user name and password is checked.

Parameters

These parameters are displayed:

- ◆ **Authentication Sequence** – Select the authentication, or authentication sequence required:
 - **Local** – User authentication is performed only locally by the switch.
 - **RADIUS** – User authentication is performed using a RADIUS server only.
 - **TACACS** – User authentication is performed using a TACACS+ server only.
 - [authentication sequence] – User authentication is performed by up to three authentication methods in the indicated sequence.

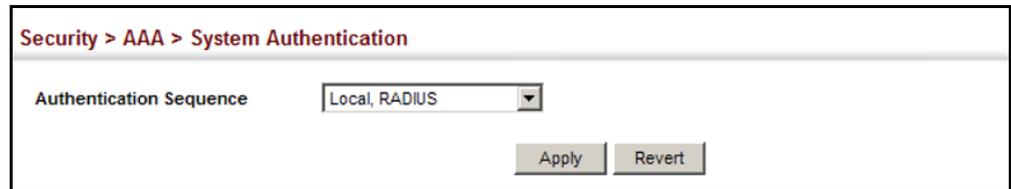
Web Interface

To configure the method(s) of controlling management access:

1. Click Security, AAA, System Authentication.
2. Specify the authentication sequence (i.e., one to three methods).

3. Click Apply.

Figure 134: Configuring the Authentication Sequence

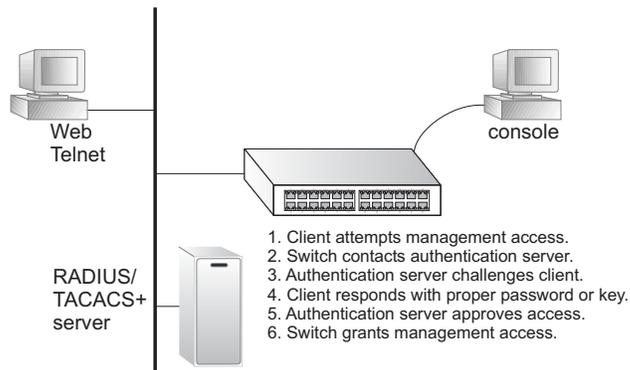


Configuring Remote Logon Authentication Servers

Use the Security > AAA > Server page to configure the message exchange parameters for RADIUS or TACACS+ remote access authentication servers.

Remote Authentication Dial-in User Service (RADIUS) and Terminal Access Controller Access Control System Plus (TACACS+) are logon authentication protocols that use software running on a central server to control access to RADIUS-aware or TACACS-aware devices on the network. An authentication server contains a database of multiple user name/password pairs with associated privilege levels for each user that requires management access to the switch.

Figure 135: Authentication Server Operation



RADIUS uses UDP while TACACS+ uses TCP. UDP only offers best effort delivery, while TCP offers a more reliable connection-oriented transport. Also, note that RADIUS encrypts only the password in the access-request packet from the client to the server, while TACACS+ encrypts the entire body of the packet.

Command Usage

- ◆ If a remote authentication server is used, you must specify the message exchange parameters for the remote authentication protocol. Both local and remote logon authentication control management access via the console port, web browser, or Telnet.
- ◆ RADIUS and TACACS+ logon authentication assign a specific privilege level for each user name/password pair. The user name, password, and privilege level must be configured on the authentication server. The encryption methods used for the authentication process must also be configured or negotiated

between the authentication server and logon client. This switch can pass authentication messages between the server and client that have been encrypted using MD5 (Message-Digest 5), TLS (Transport Layer Security), or TTLS (Tunneled Transport Layer Security).

Parameters

These parameters are displayed:

Configure Server

Server Type – Select RADIUS or TACACS+ server.

◆ RADIUS

- **Global** – Provides globally applicable RADIUS settings.
- **Server Index** – Specifies one of five RADIUS servers that may be configured. The switch attempts authentication using the listed sequence of servers. The process ends when a server either approves or denies access to a user.
- **Server IP Address** – Address of authentication server.
(A Server Index entry must be selected to display this item.)
- **Authentication Server UDP Port** – Network (UDP) port on authentication server used for authentication messages. (Range: 1-65535; Default: 1812)
- **Authentication Timeout** – The number of seconds the switch waits for a reply from the RADIUS server before it resends the request. (Range: 1-60; Default: 5)
- **Authentication Retries** – Number of times the switch tries to authenticate logon access via the authentication server. (Range: 1-5; Default: 2)
- **Set Key** – Mark this box to set or modify the encryption key.
- **Authentication Key** – Encryption key used to authenticate logon access for client. Enclose any string containing blank spaces in double quotes. (Maximum length: 48 characters)
- **Confirm Authentication Key** – Re-type the string entered in the previous field to ensure no errors were made. The switch will not change the encryption key if these two fields do not match.

◆ TACACS+

- **Global** – Provides globally applicable TACACS+ settings.
- **Server Index** – Specifies the index number of the server to be configured. The switch currently supports only one TACACS+ server.

- **Server IP Address** – Address of the TACACS+ server.
(A Server Index entry must be selected to display this item.)
- **Authentication Server TCP Port** – Network (TCP) port of TACACS+ server used for authentication messages. (Range: 1-65535; Default: 49)
- **Authentication Timeout** – The number of seconds the switch waits for a reply from the TACACS+ server before it resends the request. (Range: 1-60; Default: 5)
- **Authentication Retries** – Number of times the switch tries to authenticate logon access via the authentication server. (Range: 1-5; Default: 2)
- **Set Key** – Mark this box to set or modify the encryption key.
- **Authentication Key** – Encryption key used to authenticate logon access for client. Enclose any string containing blank spaces in double quotes. (Maximum length: 48 characters)
- **Confirm Authentication Key** – Re-type the string entered in the previous field to ensure no errors were made. The switch will not change the encryption key if these two fields do not match.

Web Interface

To configure the parameters for RADIUS or TACACS+ authentication:

1. Click Security, AAA, Server.
2. Select Configure Server from the Step list.
3. Select RADIUS or TACACS+ server type.
4. Select Global to specify the parameters that apply globally to all specified servers, or select a specific Server Index to specify the parameters that apply to a specific server.
5. To set or modify the authentication key, mark the Set Key box, enter the key, and then confirm it
6. Click Apply.

Figure 136: Configuring Remote Authentication Server (RADIUS)

The screenshot shows the configuration page for a RADIUS server. The breadcrumb is "Security > AAA > Server". Under "Server Type", "RADIUS" is selected with a radio button, and "TACACS+" is unselected. Below this, "Global" is unselected and "Server Index" has radio buttons for 1, 2, 3, 4, and 5, with "1" selected. The "Server IP Address" field contains "10.1.1.1". The "Authentication Server UDP Port (1-65535)" field contains "1813". The "Authentication Timeout (1-60)" field contains "10" with "sec" to its right. The "Authentication Retries (1-5)" field contains "5". There is a checked checkbox for "Set Key". Below it, the "Authentication Key" and "Confirm Authentication Key" fields both contain masked text "****". At the bottom right are "Apply" and "Revert" buttons.

Figure 137: Configuring Remote Authentication Server (TACACS+)

The screenshot shows the configuration page for a TACACS+ server. The breadcrumb is "Security > AAA > Server". Under "Server Type", "RADIUS" is unselected and "TACACS+" is selected with a radio button. Below this, "Global" is unselected and "Server Index" has radio buttons for 1, 2, 3, 4, and 5, with "1" selected. The "Server IP Address" field contains "10.20.30.40". The "Authentication Server TCP Port (1-65535)" field contains "200". The "Authentication Timeout (1-60)" field contains "10" with "sec" to its right. The "Authentication Retries (1-5)" field contains "5". There is a checked checkbox for "Set Key". Below it, the "Authentication Key" and "Confirm Authentication Key" fields both contain masked text "****". At the bottom right are "Apply" and "Revert" buttons.

Configuring User Accounts

Use the Security > User Accounts page to control management access to the switch based on manually configured user names and passwords.

Command Usage

- ◆ The default guest name is "guest" with the password "guest." The default administrator name is "admin" with the password "admin."
- ◆ The guest only has read access for most configuration parameters. However, the administrator has write access for all parameters governing the onboard agent. You should therefore assign a new administrator password as soon as possible, and store it in a safe place.

Parameters

These parameters are displayed:

- ◆ **User Name** – The name of the user.
(Maximum length: 32 characters; maximum number of users: 16)
- ◆ **Access Level** – Specifies the user level. (Options: 0 - Normal, 15 - Privileged)
Normal privilege level provides access to a limited number of the commands which display the current status of the switch, as well as several database clear and reset functions. Privileged level provides full access to all commands.
- ◆ **Password Type** – Specifies the following options:
 - **No Password** – No password is required for this user to log in.
 - **Plain Password** – Plain text unencrypted password.
 - **Encrypted Password** – Encrypted password.
The encrypted password is required for compatibility with legacy password settings (i.e., plain text or encrypted) when reading the configuration file during system bootup or when downloading the configuration file from a TFTP or FTP server. There is no need for you to manually configure encrypted passwords.
- ◆ **Password** – Specifies the user password. (Range: 0-32 characters, case sensitive)
- ◆ **Confirm Password** – Re-type the string entered in the previous field to ensure no errors were made. The switch will not change the password if these two fields do not match.

Web Interface

To configure user accounts:

1. Click Security, User Accounts.
2. Select Add from the Action list.
3. Specify a user name, select the user's access level, then enter a password if required and confirm it.
4. Click Apply.

Figure 138: Configuring User Accounts

Security > User Accounts

Action: Add

User Name: bob

Access Level: 15 (Privileged)

Password Type: Plain Password

Password:

Confirm Password:

Apply Revert

To show user accounts:

1. Click Security, User Accounts.
2. Select Show from the Action list.

Figure 139: Showing User Accounts

Security > User Accounts

Action: Show

User Account List Total: 3

<input type="checkbox"/>	User Name	Access Level
<input type="checkbox"/>	admin	15
<input type="checkbox"/>	guest	0
<input type="checkbox"/>	bob	15

Delete Revert

Web Authentication

Web authentication allows stations to authenticate and access the network in situations where 802.1X or Network Access authentication are infeasible or impractical. The web authentication feature allows unauthenticated hosts to request and receive a DHCP assigned IP address and perform DNS queries. All other traffic, except for HTTP protocol traffic, is blocked. The switch intercepts HTTP protocol traffic and redirects it to a switch-generated web page that facilitates user name and password authentication via RADIUS. Once authentication is successful, the web browser is forwarded on to the originally requested web page. Successful authentication is valid for all hosts connected to the port.



Note: RADIUS authentication must be activated and configured properly for the web authentication feature to work properly. (See [“Configuring Local/Remote Logon Authentication” on page 237.](#))

Note: Web authentication cannot be configured on trunk ports.

Configuring Global Settings for Web Authentication

Use the Security > Web Authentication (Configure Global) page to edit the global parameters for web authentication.

Parameters

These parameters are displayed:

- ◆ **Web Authentication Status** – Enables web authentication for the switch. (Default: Disabled)

Note that this feature must also be enabled for any port where required under the Configure Interface menu.
- ◆ **Session Timeout** – Configures how long an authenticated session stays active before it must re-authenticate itself. (Range: 300-3600 seconds; Default: 3600 seconds)
- ◆ **Quiet Period** – Configures how long a host must wait to attempt authentication again after it has exceeded the maximum allowable failed login attempts. (Range: 1-180 seconds; Default: 60 seconds)
- ◆ **Login Attempts** – Configures the amount of times a supplicant may attempt and fail authentication before it must wait the configured quiet period. (Range: 1-3 attempts; Default: 3 attempts)

Web Interface

To configure global parameters for web authentication:

1. Click Security, Web Authentication.
2. Select Configure Global from the Step list.
3. Enable web authentication globally on the switch, and adjust any of the protocol parameters as required.
4. Click Apply.

Figure 140: Configuring Global Settings for Web Authentication

Security > Web Authentication

Step: 1. Configure Global ▼

Web Authentication Status Enabled

Session Timeout (300-3600) sec

Quiet Period (1-180) sec

Login Attempts (1-3)

Configuring Interface Settings for Web Authentication

Use the Security > Web Authentication (Configure Interface) page to enable web authentication on a port, and display information for any connected hosts.

Parameters

These parameters are displayed:

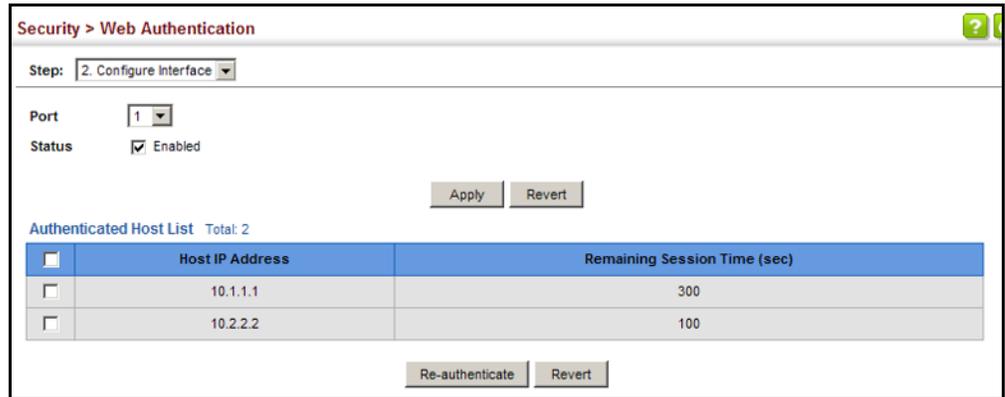
- ◆ **Port** – Indicates the port being configured.
- ◆ **Status** – Configures the web authentication status for the port.
- ◆ **Host IP Address** – Indicates the IP address of each connected host.
- ◆ **Remaining Session Time** – Indicates the remaining time until the current authorization session for the host expires.
- ◆ **Apply** – Enables web authentication if the Status box is checked.
- ◆ **Revert** – Restores the previous configuration settings.
- ◆ **Re-authenticate** – Ends all authenticated web sessions for selected host IP addresses in the Authenticated Host List, and forces the users to re-authenticate.

Web Interface

To enable web authentication for a port:

1. Click Security, Web Authentication.
2. Select Configure Interface from the Step list.
3. Set the status box to enabled for any port that requires web authentication, and click Apply
4. Mark the check box for any host addresses that need to be re-authenticated, and click Re-authenticate.

Figure 141: Configuring Interface Settings for Web Authentication



Network Access (MAC Address Authentication)

Some devices connected to switch ports may not be able to support 802.1X authentication due to hardware or software limitations. This is often true for devices such as network printers, IP phones, and some wireless access points. The switch enables network access from these devices to be controlled by authenticating device MAC addresses with a central RADIUS server.



Note: RADIUS authentication must be activated and configured properly for the MAC Address authentication feature to work properly. (See [“Configuring Remote Logon Authentication Servers”](#) on page 238.)

Note: MAC authentication cannot be configured on trunk ports.

Command Usage

- ◆ MAC address authentication controls access to the network by authenticating the MAC address of each host that attempts to connect to a switch port. Traffic received from a specific MAC address is forwarded by the switch only if the source MAC address is successfully authenticated by a central RADIUS server. While authentication for a MAC address is in progress, all traffic is blocked until authentication is completed. On successful authentication, the RADIUS server may optionally assign VLAN and quality of service settings for the switch port.
- ◆ When enabled on a port, the authentication process sends a Password Authentication Protocol (PAP) request to a configured RADIUS server. The user name and password are both equal to the MAC address being authenticated. On the RADIUS server, PAP user name and passwords must be configured in the MAC address format XX-XX-XX-XX-XX-XX (all in upper case).
- ◆ Authenticated MAC addresses are stored as dynamic entries in the switch secure MAC address table and are removed when the aging time expires. The maximum number of secure MAC addresses supported for the switch system is 1024.

- ◆ Configured static MAC addresses are added to the secure address table when seen on a switch port. Static addresses are treated as authenticated without sending a request to a RADIUS server.
- ◆ When port status changes to down, all MAC addresses mapped to that port are cleared from the secure MAC address table. Static VLAN assignments are not restored.
- ◆ The RADIUS server may optionally return a VLAN identifier list to be applied to the switch port. The following attributes need to be configured on the RADIUS server.
 - **Tunnel-Type** = VLAN
 - **Tunnel-Medium-Type** = 802
 - **Tunnel-Private-Group-ID** = 1u,2t [VLAN ID list]

The VLAN identifier list is carried in the RADIUS “Tunnel-Private-Group-ID” attribute. The VLAN list can contain multiple VLAN identifiers in the format “1u,2t,3u” where “u” indicates an untagged VLAN and “t” a tagged VLAN.
- ◆ The RADIUS server may optionally return dynamic QoS assignments to be applied to a switch port for an authenticated user. The “Filter-ID” attribute (attribute 11) can be configured on the RADIUS server to pass the following QoS information:

Table 21: Dynamic QoS Profiles

Profile	Attribute Syntax	Example
DiffServ	service-policy-in = <i>policy-map-name</i>	service-policy-in=p1
Rate Limit	rate-limit-input = <i>rate</i>	rate-limit-input=100 (kbps)
	rate-limit-output = <i>rate</i>	rate-limit-output=200 (kbps)
802.1p	switchport-priority-default = <i>value</i>	switchport-priority-default=2
IP ACL	ip-access-group-in = <i>ip-acl-name</i>	ip-access-group-in=ipV4acl
IPv6 ACL	ipv6-access-group-in = <i>ipv6-acl-name</i>	ipv6-access-group-in=ipV6acl
MAC ACL	mac-access-group-in = <i>mac-acl-name</i>	mac-access-group-in=macAcl

- ◆ Multiple profiles can be specified in the Filter-ID attribute by using a semicolon to separate each profile.
For example, the attribute “service-policy-in=pp1;rate-limit-input=100” specifies that the diffserv profile name is “pp1,” and the ingress rate limit profile value is 100 kbps.
- ◆ If duplicate profiles are passed in the Filter-ID attribute, then only the first profile is used.
For example, if the attribute is “service-policy-in=p1;service-policy-in=p2”, then the switch applies only the DiffServ profile “p1.”

- ◆ Any unsupported profiles in the Filter-ID attribute are ignored.
For example, if the attribute is “map-ip-dscp=2:3;service-policy-in=p1,” then the switch ignores the “map-ip-dscp” profile.
- ◆ When authentication is successful, the dynamic QoS information may not be passed from the RADIUS server due to one of the following conditions (authentication result remains unchanged):
 - The Filter-ID attribute cannot be found to carry the user profile.
 - The Filter-ID attribute is empty.
 - The Filter-ID attribute format for dynamic QoS assignment is unrecognizable (can not recognize the whole Filter-ID attribute).
- ◆ Dynamic QoS assignment fails and the authentication result changes from success to failure when the following conditions occur:
 - Illegal characters found in a profile value (for example, a non-digital character in an 802.1p profile value).
 - Failure to configure the received profiles on the authenticated port.
- ◆ When the last user logs off on a port with a dynamic QoS assignment, the switch restores the original QoS configuration for the port.
- ◆ When a user attempts to log into the network with a returned dynamic QoS profile that is different from users already logged on to the same port, the user is denied access.
- ◆ While a port has an assigned dynamic QoS profile, any manual QoS configuration changes only take effect after all users have logged off the port.

Configuring Global Settings for Network Access

MAC address authentication is configured on a per-port basis, however there are two configurable parameters that apply globally to all ports on the switch. Use the Security > Network Access (Configure Global) page to configure MAC address authentication aging and reauthentication time.

Parameters

These parameters are displayed:

- ◆ **Aging Status** – Enables aging for authenticated MAC addresses stored in the secure MAC address table. (Default: Disabled)

This parameter applies to authenticated MAC addresses configured by the MAC Address Authentication process described in this section, as well as to any secure MAC addresses authenticated by 802.1X, regardless of the 802.1X Operation Mode (Single-Host, Multi-Host, or MAC-Based authentication as described on [page 300](#)).

Authenticated MAC addresses are stored as dynamic entries in the switch’s secure MAC address table and are removed when the aging time expires.

The maximum number of secure MAC addresses supported for the switch system is 1024.

- ◆ **Reauthentication Time** – Sets the time period after which the switch removes an authenticated MAC address from the secure table. When the reauthentication time expires for a secure MAC address, it is removed from the secure MAC address table, and the switch will only perform the authentication process the next time it receives the MAC address packet. (Range: 120-1000000 seconds; Default: 1800 seconds)

Web Interface

To configure aging status and reauthentication time for MAC address authentication:

1. Click Security, Network Access.
2. Select Configure Global from the Step list.
3. Enable or disable aging for secure addresses, and modify the reauthentication time as required.
4. Click Apply.

Figure 142: Configuring Global Settings for Network Access

The screenshot shows the 'Security > Network Access' configuration page. At the top, there is a breadcrumb 'Security > Network Access' and a 'Step:' dropdown menu set to '1. Configure Global'. Below this, the 'Aging Status' is set to 'Enabled' with a checked checkbox. The 'Reauthentication Time (120-1000000)' is set to '30000' seconds. At the bottom right, there are 'Apply' and 'Revert' buttons.

Configuring Network Access for Ports

Use the Security > Network Access (Configure Interface - General) page to configure MAC authentication on switch ports, including enabling address authentication, setting the maximum MAC count, and enabling dynamic VLAN or dynamic QoS assignments.

Parameters

These parameters are displayed:

- ◆ **MAC Authentication**
 - **Status** – Enables MAC authentication on a port. (Default: Disabled)
 - **Intrusion** – Sets the port response to a host MAC authentication failure to either block access to the port or to pass traffic through. (Options: Block, Pass; Default: Block)

- **Max MAC Count**⁶ – Sets the maximum number of MAC addresses that can be authenticated on a port via MAC authentication; that is, the Network Access process described in this section. (Range: 1-1024; Default: 1024)
- ◆ **Network Access Max MAC Count**⁶ – Sets the maximum number of MAC addresses that can be authenticated on a port interface via all forms of authentication (including Network Access and IEEE 802.1X). (Range: 1-1024; Default: 1024)
- ◆ **Guest VLAN** – Specifies the VLAN to be assigned to the port when 802.1X Authentication or MAC authentication fails. (Range: 0-4094, where 0 means disabled; Default: Disabled)

The VLAN must already be created and active (see [“Configuring VLAN Groups” on page 147](#)). Also, when used with 802.1X authentication, intrusion action must be set for “Guest VLAN” (see [“Configuring Port Authenticator Settings for 802.1X” on page 300](#)).

A port can only be assigned to the guest VLAN in case of failed authentication if switchport mode is set to Hybrid. (See [“Adding Static Members to VLANs” on page 150](#).)

- ◆ **Dynamic VLAN** – Enables dynamic VLAN assignment for an authenticated port. When enabled, any VLAN identifiers returned by the RADIUS server through the 802.1X authentication process are applied to the port, providing the VLANs have already been created on the switch. (GVRP is not used to create the VLANs.) (Default: Enabled)

The VLAN settings specified by the first authenticated MAC address are implemented for a port. Other authenticated MAC addresses on the port must have the same VLAN configuration, or they are treated as authentication failures.

If dynamic VLAN assignment is enabled on a port and the RADIUS server returns no VLAN configuration (to the 802.1X authentication process), the authentication is still treated as a success, and the host is assigned to the default untagged VLAN.

When the dynamic VLAN assignment status is changed on a port, all authenticated addresses mapped to that port are cleared from the secure MAC address table.

- ◆ **Dynamic QoS** – Enables dynamic QoS assignment for an authenticated port. (Default: Disabled)
- ◆ **MAC Filter ID** – Allows a MAC Filter to be assigned to the port. MAC addresses or MAC address ranges present in a selected MAC Filter are exempt from authentication on the specified port (as described under [“Configuring a MAC Address Filter”](#)). (Range: 1-64; Default: None)

6. The maximum number of MAC addresses per port is 1024, and the maximum number of secure MAC addresses supported for the switch system is 1024. When the limit is reached, all new MAC addresses are treated as authentication failures.

Web Interface

To configure MAC authentication on switch ports:

1. Click Security, Network Access.
2. Select Configure Interface from the Step list.
3. Click the General button.
4. Make any configuration changes required to enable address authentication on a port, set the maximum number of secure addresses supported, the guest VLAN to use when MAC Authentication or 802.1X Authentication fails, and the dynamic VLAN and QoS assignments.
5. Click Apply.

Figure 143: Configuring Interface Settings for Network Access

Port	MAC Authentication			Network Access Max MAC Count (1-1024)	Guest VLAN (0-4094, 0: Disabled)	Dynamic VLAN	Dynamic QoS	MAC Filter ID (1-64)
	Status	Intrusion	Max MAC Count (1-1024)					
1	<input type="checkbox"/> Enabled	Block ▼	1024	1024	0	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/>
2	<input checked="" type="checkbox"/> Enabled	Pass ▼	1024	1024	0	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/>
3	<input type="checkbox"/> Enabled	Block ▼	1024	1024	0	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/>
4	<input type="checkbox"/> Enabled	Block ▼	1024	1024	0	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/>
5	<input type="checkbox"/> Enabled	Block ▼	1024	1024	0	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/>

Configuring Port Link Detection Use the Security > Network Access (Configure Interface - Link Detection) page to send an SNMP trap and/or shut down a port when a link event occurs.

Parameters

These parameters are displayed:

- ◆ **Link Detection Status** – Configures whether Link Detection is enabled or disabled for a port.
- ◆ **Condition** – The link event type which will trigger the port action.
 - **Link up** – Only link up events will trigger the port action.
 - **Link down** – Only link down events will trigger the port action.
 - **Link up and down** – All link up and link down events will trigger the port action.
- ◆ **Action** – The switch can respond in three ways to a link up or down trigger event.

- **Trap** – An SNMP trap is sent.
- **Trap and shutdown** – An SNMP trap is sent and the port is shut down.
- **Shutdown** – The port is shut down.

Web Interface

To configure link detection on switch ports:

1. Click Security, Network Access.
2. Select Configure Interface from the Step list.
3. Click the Link Detection button.
4. Modify the link detection status, trigger condition, and the response for any port.
5. Click Apply.

Figure 144: Configuring Link Detection for Network Access

Port	Link Detection Status	Condition	Action
1	<input type="checkbox"/> Enabled	Link down	Trap
2	<input type="checkbox"/> Enabled	Link up and down	Trap
3	<input type="checkbox"/> Enabled	Link down	Trap
4	<input type="checkbox"/> Enabled	Link down	Trap
5	<input type="checkbox"/> Enabled	Link down	Trap

Configuring a MAC Address Filter

Use the Security > Network Access (Configure MAC Filter) page to designate specific MAC addresses or MAC address ranges as exempt from authentication. MAC addresses present in MAC Filter tables activated on a port are treated as pre-authenticated on that port.

Command Usage

- ◆ Specified MAC addresses are exempt from authentication.
- ◆ Up to 65 filter tables can be defined.
- ◆ There is no limitation on the number of entries used in a filter table.

Parameters

These parameters are displayed:

- ◆ **Filter ID** – Adds a filter rule for the specified filter. (Range: 1-64)
- ◆ **MAC Address** – The filter rule will check ingress packets against the entered MAC address or range of MAC addresses (as defined by the MAC Address Mask).
- ◆ **MAC Address Mask** – The filter rule will check for the range of MAC addresses defined by the MAC bit mask. If you omit the mask, the system will assign the default mask of an exact match. (Range: 000000000000 - FFFFFFFF; Default: FFFFFFFF)

Web Interface

To add a MAC address filter for MAC authentication:

1. Click Security, Network Access.
2. Select Configure MAC Filter from the Step list.
3. Select Add from the Action list.
4. Enter a filter ID, MAC address, and optional mask.
5. Click Apply.

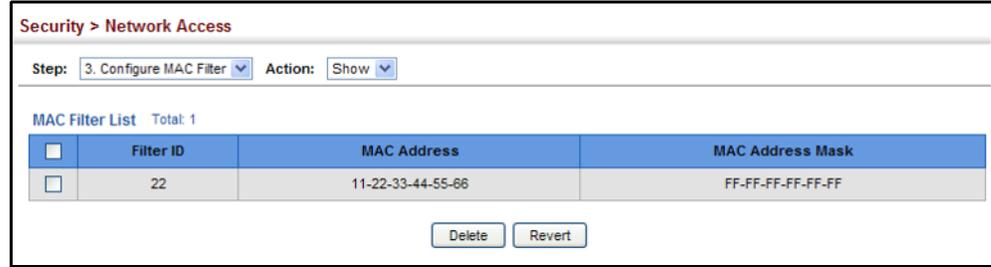
Figure 145: Configuring a MAC Address Filter for Network Access

The screenshot shows a web interface for configuring a MAC address filter. At the top, it says "Security > Network Access". Below that, there are two dropdown menus: "Step: 3. Configure MAC Filter" and "Action: Add". The main configuration area has three input fields: "Filter ID (1-64)" with the value "22", "MAC Address" with the value "11-22-33-44-55-66", and "MAC Address Mask" with the value "FFFFFFFF". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

To show the MAC address filter table for MAC authentication:

1. Click Security, Network Access.
2. Select Configure MAC Filter from the Step list.
3. Select Show from the Action list.

Figure 146: Showing the MAC Address Filter Table for Network Access



The screenshot shows the 'Security > Network Access' configuration page. At the top, there is a breadcrumb 'Security > Network Access'. Below it, a 'Step:' dropdown is set to '3. Configure MAC Filter' and an 'Action:' dropdown is set to 'Show'. The main content area is titled 'MAC Filter List' with a 'Total: 1' indicator. It contains a table with the following data:

<input type="checkbox"/>	Filter ID	MAC Address	MAC Address Mask
<input type="checkbox"/>	22	11-22-33-44-55-66	FF-FF-FF-FF-FF-FF

Below the table are two buttons: 'Delete' and 'Revert'.

Displaying Secure MAC Address Information

Use the Security > Network Access (Show Information) page to display the authenticated MAC addresses stored in the secure MAC address table. Information on the secure MAC entries can be displayed and selected entries can be removed from the table.

Parameters

These parameters are displayed:

- ◆ **Query By** – Specifies parameters to use in the MAC address query.
 - **Sort Key** – Sorts the information displayed based on MAC address, port interface, or attribute.
 - **MAC Address** – Specifies a specific MAC address.
 - **Interface** – Specifies a port interface.
 - **Attribute** – Displays static or dynamic addresses.
- ◆ **Authenticated MAC Address List**
 - **MAC Address** – The authenticated MAC address.
 - **Interface** – The port interface associated with a secure MAC address.
 - **RADIUS Server** – The IP address of the RADIUS server that authenticated the MAC address.
 - **Time** – The time when the MAC address was last authenticated.
 - **Attribute** – Indicates a static or dynamic address.

Web Interface

To display the authenticated MAC addresses stored in the secure MAC address table:

1. Click Security, Network Access.
2. Select Show Information from the Step list.
3. Use the sort key to display addresses based MAC address, interface, or attribute.

4. Restrict the displayed addresses by entering a specific address in the MAC Address field, specifying a port in the Interface field, or setting the address type to static or dynamic in the Attribute field.
5. Click Query.

Figure 147: Showing Addresses Authenticated for Network Access

The screenshot shows the 'Security > Network Access' configuration page. The 'Step' is set to '4. Show Information'. Under 'Query by:', the 'Sort Key' is 'MAC Address'. There are three checkboxes: 'MAC Address', 'Interface', and 'Attribute'. The 'Interface' dropdown is set to '1' and the 'Attribute' dropdown is set to 'Static'. A 'Query' button is visible. Below the query options, there is a table titled 'Authenticated MAC Address List' with a total of 8 entries. The table has columns for MAC Address, Interface, RADIUS Server, Time, and Attribute.

<input type="checkbox"/>	MAC Address	Interface	RADIUS Server	Time	Attribute
<input type="checkbox"/>	00-00-86-45-F2-23	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 16m 12s	Dynamic
<input type="checkbox"/>	00-00-E8-5E-E1-DD	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 32m 24s	Dynamic
<input type="checkbox"/>	00-00-E8-81-93-30	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 40m 32s	Dynamic
<input type="checkbox"/>	00-01-80-31-B8-30	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 18m 51s	Dynamic
<input type="checkbox"/>	00-01-80-36-95-D8	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 32m 22s	Dynamic

Configuring HTTPS

You can configure the switch to enable the Secure Hypertext Transfer Protocol (HTTPS) over the Secure Socket Layer (SSL), providing secure access (i.e., an encrypted connection) to the switch's web interface.

Configuring Global Settings for HTTPS

Use the Security > HTTPS (Configure Global) page to enable or disable HTTPS and specify the TCP port used for this service.

Command Usage

- ◆ Both the HTTP and HTTPS service can be enabled independently on the switch. However, you cannot configure both services to use the same TCP port. (HTTP can only be configured through the CLI using the "ip http server" command described in the *CLI Reference Guide*.)
- ◆ If you enable HTTPS, you must indicate this in the URL that you specify in your browser: `https://device[:port_number]`
- ◆ When you start HTTPS, the connection is established in this way:
 - The client authenticates the server using the server's digital certificate.

- The client and server negotiate a set of security protocols to use for the connection.
- The client and server generate session keys for encrypting and decrypting data.
- ◆ The client and server establish a secure encrypted connection.

A padlock icon should appear in the status bar for Internet Explorer 11, Mozilla Firefox 40, or Google Chrome 45, or more recent versions.

- ◆ The following web browsers and operating systems currently support HTTPS:

Table 22: HTTPS System Support

Web Browser	Operating System
Internet Explorer 11.x or later	Windows 7, 8, 10
Mozilla Firefox 40 or later	Windows 7, 8, 10, Linux
Google Chrome 45 or later	Windows 7, 8, 10

- ◆ To specify a secure-site certificate, see [“Replacing the Default Secure-site Certificate” on page 257](#).



Note: Connection to the web interface is not supported for HTTPS using an IPv6 link local address.

Parameters

These parameters are displayed:

- ◆ **HTTPS Status** – Allows you to enable/disable the HTTPS server feature on the switch. (Default: Enabled)
- ◆ **HTTPS Port** – Specifies the TCP port number used for HTTPS connection to the switch’s web interface. (Range: 1-65535, except for the following reserved ports: 1 and 25 - Linux kernel, 23 - Telnet, 80 - HTTP; Default: Port 443)

Web Interface

To configure HTTPS:

1. Click Security, HTTPS.
2. Select Configure Global from the Step list.
3. Enable HTTPS and specify the port number if required.
4. Click Apply.

Figure 148: Configuring HTTPS

Replacing the Default Secure-site Certificate

Use the Security > HTTPS (Copy Certificate) page to replace the default secure-site certificate.

When you log onto the web interface using HTTPS (for secure access), a Secure Sockets Layer (SSL) certificate appears for the switch. By default, the certificate that the web browser displays will be associated with a warning that the site is not recognized as a secure site. This is because the certificate has not been signed by an approved certification authority. If you want this warning to be replaced by a message confirming that the connection to the switch is secure, you must obtain a unique certificate and a private key and password from a recognized certification authority.



Caution: For maximum security, we recommend you obtain a unique Secure Sockets Layer certificate at the earliest opportunity. This is because the default certificate for the switch is not unique to the hardware you have purchased.

When you have obtained these, place them on your TFTP server and transfer them to the switch to replace the default (unrecognized) certificate with an authorized one.



Note: The switch must be reset for the new certificate to be activated. To reset the switch, see [“Resetting the System” on page 98](#) or type “reload” at the command prompt: `Console#reload`

Parameters

These parameters are displayed:

- ◆ **TFTP Server IP Address** – IP address of TFTP server which contains the certificate file.
- ◆ **Certificate Source File Name** – Name of certificate file stored on the TFTP server.
- ◆ **Private Key Source File Name** – Name of private key file stored on the TFTP server.

- ◆ **Private Password** – Password stored in the private key file. This password is used to verify authorization for certificate use, and is verified when downloading the certificate to the switch.
- ◆ **Confirm Password** – Re-type the string entered in the previous field to ensure no errors were made. The switch will not download the certificate if these two fields do not match.

Web Interface

To replace the default secure-site certificate:

1. Click Security, HTTPS.
2. Select Copy Certificate from the Step list.
3. Fill in the TFTP server, certificate and private key file name, and private password.
4. Click Apply.

Figure 149: Downloading the Secure-Site Certificate

The screenshot shows a web interface for configuring HTTPS. The breadcrumb is "Security > HTTPS". The "Action" dropdown is set to "Copy Certificate". Below this are five input fields: "TFTP Server IP Address" (192.168.0.4), "Certificate Source File Name" (postal-site-certificate), "Private Key Source File Name" (postal-private-key), "Private Password" (masked with dots), and "Confirm Password" (masked with dots). At the bottom right are "Apply" and "Revert" buttons.

Configuring the Secure Shell

The Berkeley-standard includes remote access tools originally designed for Unix systems. Some of these tools have also been implemented for Microsoft Windows and other environments. These tools, including commands such as *rlogin* (remote login), *rsh* (remote shell), and *rcp* (remote copy), are not secure from hostile attacks.

Secure Shell (SSH) includes server/client applications intended as a secure replacement for the older Berkeley remote access tools. SSH can also provide remote management access to this switch as a secure replacement for Telnet. When the client contacts the switch via the SSH protocol, the switch generates a public-key that the client uses along with a local user name and password for access authentication. SSH also encrypts all data transfers passing between the switch and SSH-enabled management station clients, and ensures that data traveling over the network arrives unaltered.



Note: You need to install an SSH client on the management station to access the switch for management via the SSH protocol.

Note: The switch supports both SSH Version 1.5 and 2.0 clients.

Command Usage

The SSH server on this switch supports both password and public key authentication. If password authentication is specified by the SSH client, then the password can be authenticated either locally or via a RADIUS or TACACS+ remote authentication server, as specified on the System Authentication page (page 237). If public key authentication is specified by the client, then you must configure authentication keys on both the client and the switch as described in the following section. Note that regardless of whether you use public key or password authentication, you still have to generate authentication keys on the switch (SSH Host Key Settings) and enable the SSH server (Authentication Settings).

To use the SSH server, complete these steps:

1. *Generate a Host Key Pair* – On the SSH Host Key Settings page, create a host public/private key pair.
2. *Provide Host Public Key to Clients* – Many SSH client programs automatically import the host public key during the initial connection setup with the switch. Otherwise, you need to manually create a known hosts file on the management station and place the host public key in it. An entry for a public key in the known hosts file would appear similar to the following example:

```
10.1.0.54 1024 35 15684995401867669259333946775054617325313674890836547254
15020245593199868544358361651999923329781766065830956 10825913212890233
76546801726272571413428762941301196195566782
595664104869574278881462065194174677298486546861571773939016477935594230357741
309802273708779454524083971752646358058176716709574804776117
```

3. *Import Client's Public Key to the Switch* – See [“Importing User Public Keys” on page 264](#), or use the `copy tftp public-key` command (see the `copy` command in the *CLI Reference Guide*) to copy a file containing the public key for all the SSH client's granted management access to the switch. (Note that these clients must be configured locally on the switch via the User Accounts page as described on [page 241](#).) The clients are subsequently authenticated using these keys. The current firmware only accepts public key files based on standard UNIX format as shown in the following example for an RSA Version 1 key:

```
1024 35
134108168560989392104094492015542534763164192187295892114317388005553616163105
177594083868631109291232226828519254374603100937187721199696317813662774141689
851320491172048303392543241016379975923714490119380060902539484084827178194372
288402533115952134861022902978982721353267131629432532818915045306393916643
steve@192.168.1.19
```

4. *Set the Optional Parameters* – On the SSH Settings page, configure the optional parameters, including the authentication timeout, the number of retries, and the server key size.
5. *Enable SSH Service* – On the SSH Settings page, enable the SSH server on the switch.
6. *Authentication* – One of the following authentication methods is employed:
Password Authentication (for SSH v1.5 or V2 Clients)
 - a. The client sends its password to the server.
 - b. The switch compares the client's password to those stored in memory.
 - c. If a match is found, the connection is allowed.



Note: To use SSH with only password authentication, the host public key must still be given to the client, either during initial connection or manually entered into the known host file. However, you do not need to configure the client's keys.

Public Key Authentication – When an SSH client attempts to contact the switch, the SSH server uses the host key pair to negotiate a session key and encryption method. Only clients that have a private key corresponding to the public keys stored on the switch can access it. The following exchanges take place during this process:

Authenticating SSH v1.5 Clients

- a. The client sends its RSA public key to the switch.
- b. The switch compares the client's public key to those stored in memory.
- c. If a match is found, the switch uses its secret key to generate a random 256-bit string as a challenge, encrypts this string with the user's public key, and sends it to the client.
- d. The client uses its private key to decrypt the challenge string, computes the MD5 checksum, and sends the checksum back to the switch.
- e. The switch compares the checksum sent from the client against that computed for the original string it sent. If the two checksums match, this means that the client's private key corresponds to an authorized public key, and the client is authenticated.

Authenticating SSH v2 Clients

- a. The client first queries the switch to determine if DSA public key authentication using a preferred algorithm is acceptable.
- b. If the specified algorithm is supported by the switch, it notifies the client to proceed with the authentication process. Otherwise, it rejects the request.
- c. The client sends a signature generated using the private key to the switch.

- d. When the server receives this message, it checks whether the supplied key is acceptable for authentication, and if so, it then checks whether the signature is correct. If both checks succeed, the client is authenticated.



Note: The SSH server supports up to eight client sessions. The maximum number of client sessions includes both current Telnet sessions and SSH sessions.

Note: The SSH server can be accessed using any configured IPv4 or IPv6 interface address on the switch.

Configuring the SSH Server

Use the Security > SSH (Configure Global) page to enable the SSH server and configure basic settings for authentication.



Note: You must generate DSA and RSA host keys before enabling the SSH server. See [“Generating the Host Key Pair” on page 262](#).

Parameters

These parameters are displayed:

- ◆ **SSH Server Status** – Allows you to enable/disable the SSH server on the switch. (Default: Disabled)
- ◆ **Version** – The Secure Shell version number. Version 2.0 is displayed, but the switch supports management access via either SSH Version 1.5 or 2.0 clients.
- ◆ **Authentication Timeout** – Specifies the time interval in seconds that the SSH server waits for a response from a client during an authentication attempt. (Range: 1-120 seconds; Default: 120 seconds)
- ◆ **Authentication Retries** – Specifies the number of authentication attempts that a client is allowed before authentication fails and the client has to restart the authentication process. (Range: 1-5 times; Default: 3)
- ◆ **Server-Key Size** – Specifies the SSH server key size. (Range: 512-896 bits; Default: 768)
 - The server key is a private key that is never shared outside the switch.
 - The host key is shared with the SSH client, and is fixed at 1024 bits.

Web Interface

To configure the SSH server:

1. Click Security, SSH.
2. Select Configure Global from the Step list.

3. Enable the SSH server.
4. Adjust the authentication parameters as required.
5. Click Apply.

Figure 150: Configuring the SSH Server

Security > SSH

Step: 1. Configure Global

SSH Server Status	<input checked="" type="checkbox"/> Enabled
Version	2.0
Authentication Timeout (1-120)	<input type="text" value="120"/> sec
Authentication Retries (1-5)	<input type="text" value="3"/>
Server-Key Size (512-896)	<input type="text" value="768"/>

Apply Revert

Generating the Host Key Pair

Use the Security > SSH (Configure Host Key - Generate) page to generate a host public/private key pair used to provide secure communications between an SSH client and the switch. After generating this key pair, you must provide the host public key to SSH clients and import the client's public key to the switch as described in the section "Importing User Public Keys" on page 264.



Note: A host key pair must be configured on the switch before you can enable the SSH server. See "Configuring the SSH Server" on page 261.

Parameters

These parameters are displayed:

Generate Host Key

- ◆ **Host-Key Type** – The key type used to generate the host key pair (i.e., public and private keys). (Range: RSA (Version 1), DSA (Version 2), Both; Default: Both)

The SSH server uses RSA or DSA for key exchange when the client first establishes a connection with the switch, and then negotiates with the client to select either DES (56-bit) or 3DES (168-bit) for data encryption.



Note: The switch uses only RSA Version 1 for SSHv1.5 clients and DSA Version 2 for SSHv2 clients.

Web Interface

To generate the SSH host key pair:

1. Click Security, SSH.
2. Select Configure Host Key from the Step list.
3. Select Generate from the Action list.
4. Select the host-key type from the drop-down box.
5. Click Apply.

Figure 151: Generating the SSH Host Key Pair

The screenshot shows a web interface for configuring SSH. At the top, it says "Security > SSH". Below that, there are two dropdown menus: "Step:" with "2. Configure Host Key" selected, and "Action:" with "Generate" selected. Below these, there is a "Host-Key Type" dropdown menu with "Both" selected. At the bottom right, there are two buttons: "Apply" and "Revert".

Showing the Host Key Pair Use the Security > SSH (Configure Host Key - Show) page to show the host public key used to provide secure communications between an SSH client and the switch.

Parameters

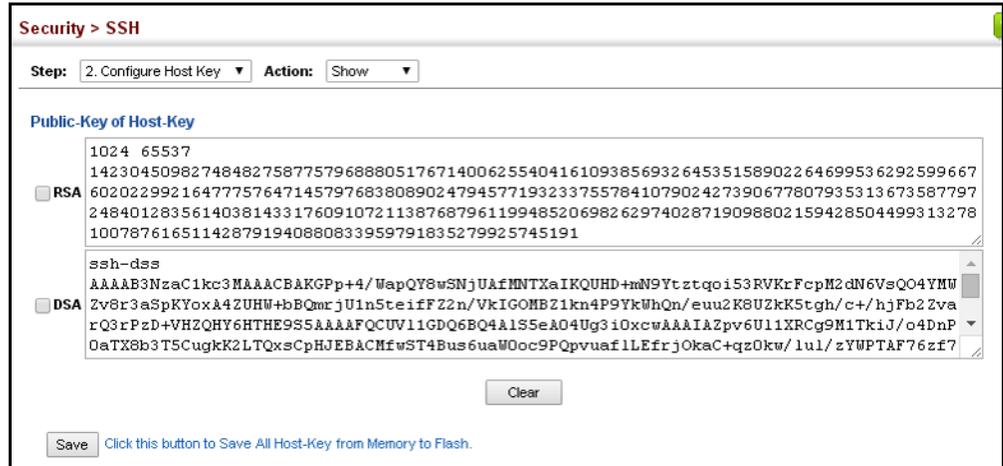
These parameters are displayed:

- ◆ **Public-Key of Host-Key** – The host public key generated by the switch.
- ◆ **Clear** – Clears the RSA or DSA public keys when the check box is selected.
- ◆ **Save** – Saves the host key from RAM (i.e., volatile memory) to flash memory. Otherwise, the host key pair is stored to RAM by default.

To display or clear the SSH host key pair:

1. Click Security, SSH.
2. Select Configure Host Key from the Step list.
3. Select Show from the Action list.
4. Select the option to save the host key from memory to flash by clicking Save, or select the host-key type to clear and click Clear.

Figure 152: Showing the SSH Host Key Pair



Importing User Public Keys

Use the Security > SSH (Configure User Key - Copy) page to upload a user's public key to the switch. This public key must be stored on the switch for the user to be able to log in using the public key authentication mechanism. If the user's public key does not exist on the switch, SSH will revert to the interactive password authentication mechanism to complete authentication.

Parameters

These parameters are displayed:

- ◆ **User Name** – This drop-down box selects the user who's public key you wish to manage. Note that you must first create users on the User Accounts page (see ["Configuring User Accounts" on page 241](#)).
- ◆ **User Key Type** – The type of public key to upload.
 - RSA: The switch accepts a RSA version 1 encrypted public key.
 - DSA: The switch accepts a DSA version 2 encrypted public key.

The SSH server uses RSA or DSA for key exchange when the client first establishes a connection with the switch, and then negotiates with the client to select either DES (56-bit) or 3DES (168-bit) for data encryption.

The switch uses only RSA Version 1 for SSHv1.5 clients and DSA Version 2 for SSHv2 clients.

- ◆ **TFTP Server IP Address** – The IP address of the TFTP server that contains the public key file you wish to import.
- ◆ **Source File Name** – The public key file to upload.

Web Interface

To copy the SSH user's public key:

1. Click Security, SSH.
2. Select Configure User Key from the Step list.
3. Select Copy from the Action list.
4. Select the user name and the public-key type from the respective drop-down boxes, input the TFTP server IP address and the public key source file name.
5. Click Apply.

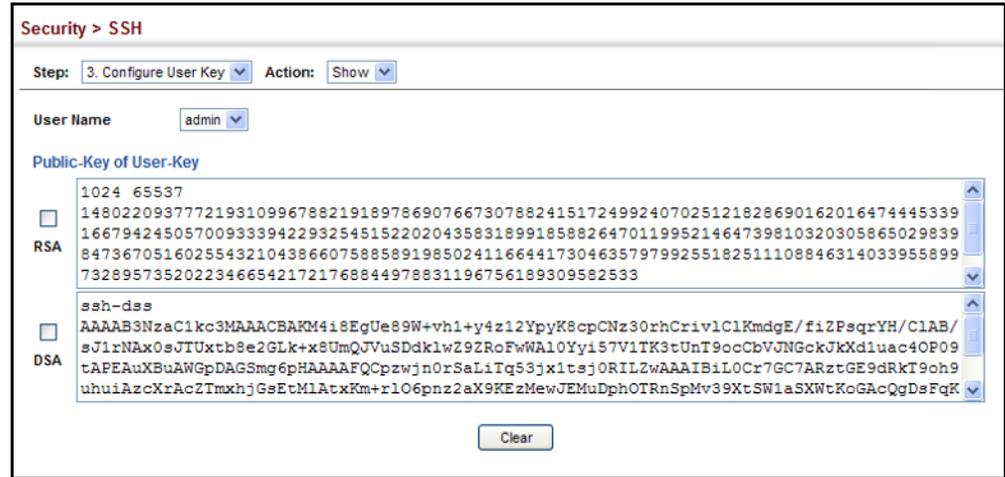
Figure 153: Copying the SSH User's Public Key

The screenshot shows a web interface for configuring SSH user keys. The breadcrumb is "Security > SSH". At the top, there are two dropdown menus: "Step:" set to "3. Configure User Key" and "Action:" set to "Copy". Below these are four input fields: "User Name" with a dropdown menu showing "steve", "User-Key Type" with a dropdown menu showing "RSA", "TFTP Server IP Address" with a text input field containing "192.168.0.61", and "Source File Name" with a text input field containing "rsa.pub". At the bottom right, there are two buttons: "Apply" and "Revert".

To display or clear the SSH user's public key:

1. Click Security, SSH.
2. Select Configure User Key from the Step list.
3. Select Show from the Action list.
4. Select a user from the User Name list.
5. Select the host-key type to clear.
6. Click Clear.

Figure 154: Showing the SSH User's Public Key



Access Control Lists

Access Control Lists (ACL) provide packet filtering for IPv4 frames (based on address, protocol, Layer 4 protocol port number or TCP control code), IPv6 frames (based on address or DSCP traffic class, DSCP, next header type, or flow label), or any frames (based on MAC address or Ethernet type). To filter incoming packets, first create an access list, add the required rules, and then bind the list to a specific port.

Configuring Access Control Lists –

An ACL is a sequential list of permit or deny conditions that apply to IP addresses, MAC addresses, or other more specific criteria. This switch tests ingress packets against the conditions in an ACL one by one. A packet will be accepted as soon as it matches a permit rule, or dropped as soon as it matches a deny rule. If no rules match, the packet is accepted.

Command Usage

The following restrictions apply to ACLs:

- ◆ The maximum number of ACLs is 256.
- ◆ The maximum number of rules per ACL is 96.
- ◆ An ACL can have up to 96 rules. However, due to resource restrictions, the average number of rules bound to the ports should not exceed 20.
- ◆ The maximum number of rules that can be bound to the ports is 64 for each of the following list types: MAC ACLs, IP ACLs (including Standard and Extended ACLs), IPv6 Standard ACLs, and IPv6 Extended ACLs.
- ◆ The maximum number of rules (Access Control Entries, or ACEs) stated above is the worst case scenario. In practice, the switch compresses the ACEs in TCAM (a

hardware table used to store ACEs), but the actual maximum number of ACEs possible depends on too many factors to be precisely determined. It depends on the amount of hardware resources reserved at runtime for this purpose.

Auto ACE Compression is a software feature used to compress all the ACEs of an ACL to utilize hardware resources more efficiently. Without compression, one ACE would occupy a fixed number of entries in TCAM. So if one ACL includes 25 ACEs, the ACL would need $(25 * n)$ entries in TCAM, where "n" is the fixed number of TCAM entries needed for one ACE. When compression is employed, before writing the ACE into TCAM, the software compresses the ACEs to reduce the number of required TCAM entries. For example, one ACL may include 128 ACEs which classify a continuous IP address range like 192.168.1.0~255. If compression is disabled, the ACL would occupy $(128*n)$ entries of TCAM, using up nearly all of the hardware resources. When using compression, the 128 ACEs are compressed into one ACE classifying the IP address as 192.168.1.0/24, which requires only "n" entries in TCAM. The above example is an ideal case for compression. The worst case would be if no any ACE can be compressed, in which case the used number of TCAM entries would be the same as without compression. It would also require more time to process the ACEs.

- ◆ If no matches are found down to the end of the list, the traffic is denied. For this reason, frequently hit entries should be placed at the top of the list. There is an implied deny for traffic that is not explicitly permitted. Also, note that a single-entry ACL with only one deny entry has the effect of denying all traffic. You should therefore use at least one permit statement in an ACL or all traffic will be blocked.

Because the switch stops testing after the first match, the order of the conditions is critical. If no conditions match, the packet will be denied.

The order in which active ACLs are checked is as follows:

1. User-defined rules in IP and MAC ACLs for ingress ports are checked in parallel.
2. Rules within an ACL are checked in the configured order, from top to bottom.
3. If the result of checking an IP ACL is to permit a packet, but the result of a MAC ACL on the same packet is to deny it, the packet will be denied (because the decision to deny a packet has a higher priority for security reasons). A packet will also be denied if the IP ACL denies it and the MAC ACL accepts it.

Showing TCAM Utilization Use the Security > ACL (Configure ACL - Show TCAM) page to show utilization parameters for TCAM (Ternary Content Addressable Memory), including the number policy control entries in use, the number of free entries, and the overall percentage of TCAM in use.

Command Usage

Policy control entries (PCEs) are used by various system functions which rely on rule-based searches, including Access Control Lists (ACLs), IP Source Guard filter rules, Quality of Service (QoS) processes, QinQ, MAC-based VLANs, VLAN translation, or traps.

For example, when binding an ACL to a port, each rule in an ACL will use two PCEs; and when setting an IP Source Guard filter rule for a port, the system will also use two PCEs.

Parameters

These parameters are displayed:

- ◆ **Total Policy Control Entries** – The number policy control entries in use.
- ◆ **Free Policy Control Entries** – The number of policy control entries available for use.
- ◆ **Entries Used by System** – The number of policy control entries used by the operating system.
- ◆ **Entries Used by User** – The number of policy control entries used by configuration settings, such as access control lists.
- ◆ **TCAM Utilization** – The overall percentage of TCAM in use.

Web Interface

To show information on TCAM utilization:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Show TCAM from the Action list.

Figure 155: Showing TCAM Utilization



Security > ACL	
Step:	2. Configure ACL
Action:	Show TCAM
Total Policy Control Entries	960
Free Policy Control Entries	960
Entries Used by System	0
Entries Used by User	0
TCAM Utilization	0.0%

Setting the ACL Name and Type Use the Security > ACL (Configure ACL - Add) page to create an ACL.

Parameters

These parameters are displayed:

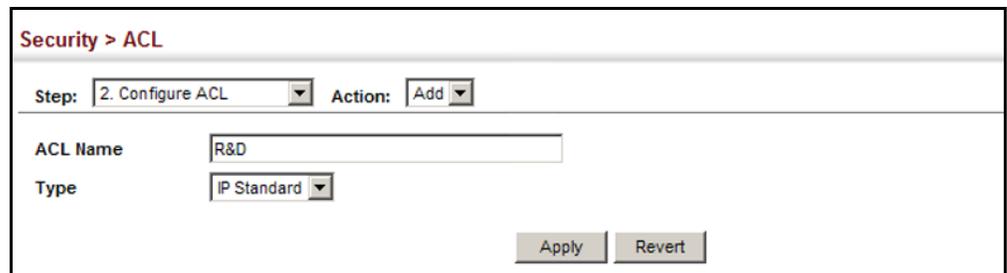
- ◆ **ACL Name** – Name of the ACL. (Maximum length: 32 characters)
- ◆ **Type** – The following filter modes are supported:
 - **IP Standard:** IPv4 ACL mode filters packets based on the source IPv4 address.
 - **IP Extended:** IPv4 ACL mode filters packets based on the source or destination IPv4 address, as well as the protocol type and protocol port number. If the “TCP” protocol is specified, then you can also filter packets based on the TCP control code.
 - **IPv6 Standard:** IPv6 ACL mode filters packets based on the source IPv6 address.
 - **IPv6 Extended:** IPv6 ACL mode filters packets based on the source or destination IP address.
 - **MAC** – MAC ACL mode filters packets based on the source or destination MAC address and the Ethernet frame type (RFC 1060).
 - **ARP** – ARP ACL specifies static IP-to-MAC address bindings used for ARP inspection (see [“ARP Inspection” on page 285](#)).

Web Interface

To configure the name and type of an ACL:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Add from the Action list.
4. Fill in the ACL Name field, and select the ACL type.
5. Click Apply.

Figure 156: Creating an ACL

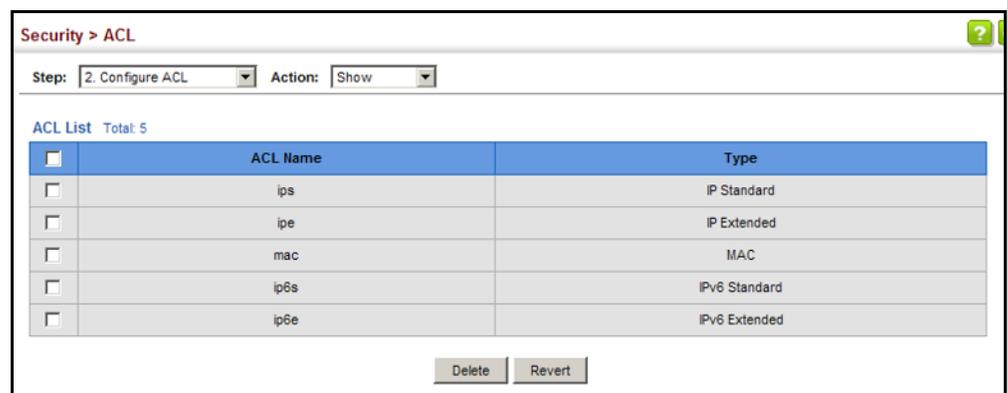


The screenshot shows the 'Security > ACL' configuration page. At the top, there is a breadcrumb 'Security > ACL'. Below it, there are two dropdown menus: 'Step: 2. Configure ACL' and 'Action: Add'. The main form has two fields: 'ACL Name' with the value 'R&D' and 'Type' with the value 'IP Standard'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

To show a list of ACLs:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Show from the Action list.

Figure 157: Showing a List of ACLs



The screenshot shows the 'Security > ACL' configuration page with the 'Action' dropdown set to 'Show'. Below the form, there is a table titled 'ACL List Total: 5'. The table has three columns: a checkbox, 'ACL Name', and 'Type'. The table lists five ACLs: 'ips' (IP Standard), 'ipe' (IP Extended), 'mac' (MAC), 'ip6s' (IPv6 Standard), and 'ip6e' (IPv6 Extended). At the bottom right, there are two buttons: 'Delete' and 'Revert'.

<input type="checkbox"/>	ACL Name	Type
<input type="checkbox"/>	ips	IP Standard
<input type="checkbox"/>	ipe	IP Extended
<input type="checkbox"/>	mac	MAC
<input type="checkbox"/>	ip6s	IPv6 Standard
<input type="checkbox"/>	ip6e	IPv6 Extended

Configuring a Standard IPv4 ACL Use the Security > ACL (Configure ACL - Add Rule - IP Standard) page to configure a Standard IPv4 ACL.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to show in the Name list.
- ◆ **Name** – Shows the names of ACLs matching the selected type.
- ◆ **Action** – An ACL can contain any combination of permit or deny rules.
- ◆ **Address Type** – Specifies the source IP address. Use “Any” to include all possible addresses, “Host” to specify a specific host address in the Address field, or “IP” to specify a range of addresses with the Address and Subnet Mask fields. (Options: Any, Host, IP; Default: Any)
- ◆ **Source IP Address** – Source IP address.
- ◆ **Source Subnet Mask** – A subnet mask containing four integers from 0 to 255, each separated by a period. The mask uses 1 bits to indicate “match” and 0 bits to indicate “ignore.” The mask is bitwise ANDed with the specified source IP address, and compared with the address for each IP packet entering the port(s) to which this ACL has been assigned.

Web Interface

To add rules to an IPv4 Standard ACL:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Add Rule from the Action list.
4. Select IP Standard from the Type list.
5. Select the name of an ACL from the Name list.
6. Specify the action (i.e., Permit or Deny).
7. Select the address type (Any, Host, or IP).
8. If you select “Host,” enter a specific address. If you select “IP,” enter a subnet address and the mask for an address range.
9. Click Apply.

Figure 158: Configuring a Standard IPv4 ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, the 'Step' is '1. Configure ACL' and the 'Action' is 'Add Rule'. The 'Type' is set to 'IP Standard' (selected with a radio button). The 'Name' is 'R&D'. The 'Action' is 'Permit'. The 'Address Type' is 'Host'. The 'Source IP Address' is '10.1.1.21' and the 'Source Subnet Mask' is '255.255.255.255'. There are 'Apply' and 'Revert' buttons at the bottom right.

Configuring an Extended IPv4 ACL Use the Security > ACL (Configure ACL - Add Rule - IP Extended) page to configure an Extended IPv4 ACL.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to show in the Name list.
- ◆ **Name** – Shows the names of ACLs matching the selected type.
- ◆ **Action** – An ACL can contain any combination of permit or deny rules.
- ◆ **Source/Destination Address Type** – Specifies the source or destination IP address type. Use “Any” to include all possible addresses, “Host” to specify a specific host address in the Address field, or “IP” to specify a range of addresses with the Address and Subnet Mask fields. (Options: Any, Host, IP; Default: Any)
- ◆ **Source/Destination IP Address** – Source or destination IP address.
- ◆ **Source/Destination Subnet Mask** – Subnet mask for source or destination address. (See the description for Subnet Mask on [page 271](#).)
- ◆ **Source/Destination Port** – Source/destination port number for the specified protocol type. (Range: 0-65535)
- ◆ **Source/Destination Port Bit Mask** – Decimal number representing the port bits to match. (Range: 0-65535)
- ◆ **Protocol** – Specifies the protocol type to match as TCP, UDP or Others, where others indicates a specific protocol number (0-255). (Options: TCP, UDP, Others; Default: Others)
- ◆ **Service Type** – Packet priority settings based on the following criteria:
 - **ToS** – Type of Service level. (Range: 0-15)

- **Precedence** – IP precedence level. (Range: 0-7)
- **DSCP** – DSCP priority level. (Range: 0-63)
- ◆ **Control Code** – Decimal number (representing a bit string) that specifies flag bits in byte 14 of the TCP header. (Range: 0-63)
- ◆ **Control Code Bit Mask** – Decimal number representing the code bits to match. (Range: 0-63)

The control bit mask is a decimal number (for an equivalent binary bit mask) that is applied to the control code. Enter a decimal number, where the equivalent binary bit “1” means to match a bit and “0” means to ignore a bit. The following bits may be specified:

- 1 (fin) – Finish
- 2 (syn) – Synchronize
- 4 (rst) – Reset
- 8 (psh) – Push
- 16 (ack) – Acknowledgement
- 32 (urg) – Urgent pointer

For example, use the code value and mask below to catch packets with the following flags set:

- SYN flag valid, use control-code 2, control bit mask 2
- Both SYN and ACK valid, use control-code 18, control bit mask 18
- SYN valid and ACK invalid, use control-code 2, control bit mask 18

Web Interface

To add rules to an IPv4 Extended ACL:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Add Rule from the Action list.
4. Select IP Extended from the Type list.
5. Select the name of an ACL from the Name list.
6. Specify the action (i.e., Permit or Deny).
7. Select the address type (Any, Host, or IP).
8. If you select “Host,” enter a specific address. If you select “IP,” enter a subnet address and the mask for an address range.

9. Set any other required criteria, such as service type, protocol type, or control code.
10. Click Apply.

Figure 159: Configuring an Extended IPv4 ACL

The screenshot shows the 'Security > ACL' configuration page. The 'Step' is '1. Configure ACL' and the 'Action' is 'Add Rule'. The 'Type' is 'IP Extended'. The 'Name' is 'R&D#2'. The 'Action' is 'Permit'. The 'Source Address Type' is 'IP', with 'Source IP Address' set to '10.7.1.0' and 'Source Subnet Mask' set to '255.255.255.0'. The 'Destination Address Type' is 'Any', with 'Destination IP Address' set to '0.0.0.0' and 'Destination Subnet Mask' set to '0.0.0.0'. The 'Protocol' is 'TCP (6)'. The 'Service Type' is 'Precedence (0-7)'. There are 'Apply' and 'Revert' buttons at the bottom.

Configuring a Standard IPv6 ACL Use the Security > ACL (Configure ACL - Add Rule - IPv6 Standard) page to configure a Standard IPv6 ACL.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to show in the Name list.
- ◆ **Name** – Shows the names of ACLs matching the selected type.
- ◆ **Action** – An ACL can contain any combination of permit or deny rules.
- ◆ **Source Address Type** – Specifies the source IP address. Use “Any” to include all possible addresses, “Host” to specify a specific host address in the Address field, or “IPv6-Prefix” to specify a range of addresses. (Options: Any, Host, IPv6-Prefix; Default: Any)
- ◆ **Source IPv6 Address** – An IPv6 source address or network class. The address must be formatted according to RFC 2373 “IPv6 Addressing Architecture,” using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields.

- ◆ **Source Prefix Length** – A decimal value indicating how many contiguous bits (from the left) of the address comprise the prefix (i.e., the network portion of the address). (Range: 0-128 bits)

Web Interface

To add rules to a Standard IPv6 ACL:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Add Rule from the Action list.
4. Select IPv6 Standard from the Type list.
5. Select the name of an ACL from the Name list.
6. Specify the action (i.e., Permit or Deny).
7. Select the source address type (Any, Host, or IPv6-prefix).
8. If you select “Host,” enter a specific address. If you select “IPv6-prefix,” enter a subnet address and the prefix length.
9. Click Apply.

Figure 160: Configuring a Standard IPv6 ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, there are two dropdown menus: 'Step: 1. Configure ACL' and 'Action: Add Rule'. Below this, the 'Type' section has five radio buttons: 'IP Standard', 'IP Extended', 'MAC', 'IPv6 Standard' (which is selected), and 'IPv6 Extended'. The 'Name' field contains 'R&D#6S'. The 'Action' dropdown is set to 'Permit'. The 'Source Address Type' dropdown is set to 'Host'. The 'Source IPv6 Address' field contains '2009:DB9:2229::79'. The 'Source Prefix Length (0-128)' field contains '128'. At the bottom right, there are 'Apply' and 'Revert' buttons.

Configuring an Extended IPv6 ACL Use the Security > ACL (Configure ACL - Add Rule - IPv6 Extended) page to configure an Extended IPv6 ACL.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to show in the Name list.
- ◆ **Name** – Shows the names of ACLs matching the selected type.
- ◆ **Action** – An ACL can contain any combination of permit or deny rules.
- ◆ **Destination Address Type** – Specifies the destination IP address type. Use “Any” to include all possible addresses, “Host” to specify a specific host address in the Address field, or “IPv6-Prefix” to specify a range of addresses. (Options: Any, Host, IPv6-Prefix; Default: Any)
- ◆ **Destination IPv6 Address** – An IPv6 address or network class. The address must be formatted according to RFC 2373 “IPv6 Addressing Architecture,” using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields.
- ◆ **Destination Prefix Length** – A decimal value indicating how many contiguous bits (from the left) of the address comprise the prefix; i.e., the network portion of the address. (Range: 0-128 bits)
- ◆ **Next Header** – Identifies the type of header immediately following the IPv6 header. (Range: 0-255)

Optional internet-layer information is encoded in separate headers that may be placed between the IPv6 header and the upper-layer header in a packet. There are a small number of such extension headers, each identified by a distinct Next Header value. IPv6 supports the values defined for the IPv4 Protocol field in RFC 1700, and includes these commonly used headers:

- 0 : Hop-by-Hop Options (RFC 2460)
- 6 : TCP Upper-layer Header (RFC 1700)
- 17 : UDP Upper-layer Header (RFC 1700)
- 43 : Routing (RFC 2460)
- 44 : Fragment (RFC 2460)
- 50 : Encapsulating Security Payload (RFC 2406)
- 51 : Authentication (RFC 2402)
- 60 : Destination Options (RFC 2460)

Web Interface

To add rules to an Extended IPv6 ACL:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.

3. Select Add Rule from the Action list.
4. Select IPv6 Extended from the Type list.
5. Select the name of an ACL from the Name list.
6. Specify the action (i.e., Permit or Deny).
7. Select the address type (Any, Host or IPv6-prefix).
8. If you select "Host," enter a specific address. If you select "IPv6-prefix," enter a subnet address and prefix length.
9. Click Apply.

Figure 161: Configuring an Extended IPv6 ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, there are two dropdown menus: 'Step: 1. Configure ACL' and 'Action: Add Rule'. Below this, the 'Type' section has five radio buttons: 'IP Standard', 'IP Extended', 'MAC', 'IPv6 Standard', and 'IPv6 Extended', with 'IPv6 Extended' selected. The 'Name' field is a dropdown menu showing 'ipv6e'. The 'Action' field is a dropdown menu showing 'Permit'. The 'Destination Address Type' field is a dropdown menu showing 'IPv6-Prefix'. The 'Destination IPv6 Address' field is a text input containing '2009:db9:2229::78'. The 'Destination Prefix Length (0-128)' field is a text input containing '8'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

Configuring a MAC ACL Use the Security > ACL (Configure ACL - Add Rule - MAC) page to configure a MAC ACL based on hardware addresses, packet format, and Ethernet type.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to show in the Name list.
- ◆ **Name** – Shows the names of ACLs matching the selected type.
- ◆ **Action** – An ACL can contain any combination of permit or deny rules.
- ◆ **Source/Destination Address Type** – Use "Any" to include all possible addresses, "Host" to indicate a specific MAC address, or "MAC" to specify an address range with the Address and Bit Mask fields. (Options: Any, Host, MAC; Default: Any)
- ◆ **Source/Destination MAC Address** – Source or destination MAC address.

- ◆ **Source/Destination Bit Mask** – Hexadecimal mask for source or destination MAC address.
- ◆ **Packet Format** – This attribute includes the following packet types:
 - **Any** – Any Ethernet packet type.
 - **Untagged-eth2** – Untagged Ethernet II packets.
 - **Untagged-802.3** – Untagged Ethernet 802.3 packets.
 - **Tagged-eth2** – Tagged Ethernet II packets.
 - **Tagged-802.3** – Tagged Ethernet 802.3 packets.
- ◆ **VID** – VLAN ID. (Range: 1-4094)
- ◆ **VID Bit Mask** – VLAN bit mask. (Range: 0-4095)
- ◆ **Ethernet Type** – This option can only be used to filter Ethernet II formatted packets. (Range: 0-ffff hex.)

A detailed listing of Ethernet protocol types can be found in RFC 1060. A few of the more common types include 0800 (IP), 0806 (ARP), 8137 (IPX).
- ◆ **Ethernet Type Bit Mask** – Protocol bit mask. (Range: 0-ffff hex.)

Web Interface

To add rules to a MAC ACL:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Add Rule from the Action list.
4. Select MAC from the Type list.
5. Select the name of an ACL from the Name list.
6. Specify the action (i.e., Permit or Deny).
7. Select the address type (Any, Host, or MAC).
8. If you select "Host," enter a specific address (e.g., 11-22-33-44-55-66). If you select "MAC," enter a base address and a hexadecimal bit mask for an address range.
9. Set any other required criteria, such as VID, Ethernet type, or packet format.
10. Click Apply.

Figure 162: Configuring a MAC ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, it indicates 'Step: 1. Configure ACL' and 'Action: Add Rule'. Below this, there are radio buttons for 'Type': IP Standard, IP Extended, MAC (selected), IPv6 Standard, and IPv6 Extended. The 'Name' field is set to 'mac'. The 'Action' is set to 'Permit'. The 'Source Address Type' is 'Any', and the 'Destination Address Type' is 'Any'. Both 'Source MAC Address' and 'Destination MAC Address' are set to '00-00-00-00-00-00'. Both 'Source Bit Mask' and 'Destination Bit Mask' are also set to '00-00-00-00-00-00'. The 'Packet Format' is set to 'Any'. The 'VID (1-4094)' is set to '12' and the 'VID Bit Mask (0-4095)' is set to '4095'. There are empty fields for 'Ethernet Type' and 'Ethernet Type Bit Mask', both with '(0000-FFFF, hexadecimal value)' as a hint. At the bottom right, there are 'Apply' and 'Revert' buttons.

Configuring an ARP ACL

Use the Security > ACL (Configure ACL - Add Rule - ARP) page to configure ACLs based on ARP message addresses. ARP Inspection can then use these ACLs to filter suspicious traffic (see [“Configuring Global Settings for ARP Inspection”](#) on page 286).

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to show in the Name list.
- ◆ **Name** – Shows the names of ACLs matching the selected type.
- ◆ **Action** – An ACL can contain any combination of permit or deny rules.
- ◆ **Packet Type** – Indicates an ARP request, ARP response, or either type. (Range: IP, Request, Response; Default: IP)
- ◆ **Source/Destination IP Address Type** – Specifies the source or destination IPv4 address. Use “Any” to include all possible addresses, “Host” to specify a specific host address in the Address field, or “IP” to specify a range of addresses with the Address and Mask fields. (Options: Any, Host, IP; Default: Any)
- ◆ **Source/Destination IP Address** – Source or destination IP address.
- ◆ **Source/Destination IP Subnet Mask** – Subnet mask for source or destination address. (See the description for Subnet Mask on [page 271](#).)

- ◆ **Source/Destination MAC Address Type** – Use “Any” to include all possible addresses, “Host” to indicate a specific MAC address, or “MAC” to specify an address range with the Address and Mask fields. (Options: Any, Host, MAC; Default: Any)
- ◆ **Source/Destination MAC Address** – Source or destination MAC address.
- ◆ **Source/Destination MAC Bit Mask** – Hexadecimal mask for source or destination MAC address.
- ◆ **Log** – Logs a packet when it matches the access control entry.

Web Interface

To add rules to an ARP ACL:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Add Rule from the Action list.
4. Select ARP from the Type list.
5. Select the name of an ACL from the Name list.
6. Specify the action (i.e., Permit or Deny).
7. Select the packet type (Request, Response, All).
8. Select the address type (Any, Host, or IP).
9. If you select “Host,” enter a specific address (e.g., 11-22-33-44-55-66). If you select “IP,” enter a base address and a hexadecimal bit mask for an address range.
10. Enable logging if required.
11. Click Apply.

Figure 163: Configuring a ARP ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, the 'Step' is '2. Configure ACL' and the 'Action' is 'Add Rule'. The 'Type' is set to 'ARP'. The 'Name' is 'R&D#7ARP'. The 'Action' is 'Permit'. The 'Packet Type' is 'All'. The 'Source IP Address Type' is 'Any', 'Source IP Address' is '0.0.0.0', and 'Source IP Subnet Mask' is '0.0.0.0'. The 'Destination IP Address Type' is 'Any', 'Destination IP Address' is '0.0.0.0', and 'Destination IP Subnet Mask' is '0.0.0.0'. The 'Source MAC Address Type' is 'Any', 'Source MAC Address' is '00-00-00-00-00-00', and 'Source MAC Bit Mask' is '00-00-00-00-00-00'. The 'Destination MAC Address Type' is 'Any', 'Destination MAC Address' is '00-00-00-00-00-00', and 'Destination MAC Bit Mask' is '00-00-00-00-00-00'. There is a 'Log' checkbox which is unchecked. At the bottom, there are 'Apply' and 'Revert' buttons.

Binding a Port to an Access Control List

After configuring ACLs, use the Security > ACL (Configure Interface – Configure) page to bind the ports that need to filter traffic to the appropriate ACLs. Only one access list (IPv4, IPv6 or MAC) can be assigned to an ingress or egress port.

Command Usage

- ◆ This switch supports ACLs for ingress filtering only.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to bind to a port.
- ◆ **Port** – Port identifier.
- ◆ **ACL** – ACL used for ingress or egress packets.
- ◆ **Counter** – Enables counter for ACL statistics.

Web Interface

To bind an ACL to a port:

1. Click Security, ACL.
2. Select Configure Interface from the Step list.
3. Select Configure from the Action list.
4. Select IP, MAC or IPv6 from the Type list.
5. Select a port.

6. Select the name of an ACL from the ACL list.
7. Click Apply.

Figure 164: Binding a Port to an ACL

The screenshot shows a web interface for configuring an ACL. At the top, it says "Security > ACL". Below that, there are two dropdown menus: "Step: 2. Configure Interface" and "Action: Configure". Under the "Type" section, there are three radio buttons: "IP" (selected), "MAC", and "IPv6". Below that is a "Port" dropdown menu with "1" selected. There are two sections: "Ingress" and "Egress". Under "Ingress", there is an "ACL" dropdown menu with "R&D" selected and a checked "Counter" checkbox. Under "Egress", there is an "ACL" dropdown menu with "R&D#2" selected and a checked "Counter" checkbox. At the bottom right, there are two buttons: "Apply" and "Revert".

Configuring ACL Mirroring

After configuring ACLs, use the Security > ACL > Configure Interface (Add Mirror) page to mirror traffic matching an ACL from one or more source ports to a target port for real-time analysis. You can then attach a logic analyzer or RMON probe to the target port and study the traffic crossing the source VLAN(s) in a completely unobtrusive manner.

Command Usage

ACL-based mirroring is only used for ingress traffic. To mirror an ACL, follow these steps:

1. Create an ACL as described in the preceding sections.
2. Add one or more mirrored ports to ACL as described under [“Binding a Port to an Access Control List” on page 281](#).
3. Use the Add Mirror page to specify the ACL and the destination port to which matching traffic will be mirrored.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier.
- ◆ **ACL** – ACL used for ingress packets.

Web Interface

To bind an ACL to a port:

1. Click Security, ACL.
2. Select Configure Interface from the Step list.
3. Select Add Mirror from the Action list.
4. Select a port.
5. Select the name of an ACL from the ACL list.
6. Click Apply.

Figure 165: Configuring ACL Mirroring

Security > ACL

Step: 3. Configure Interface Action: Add Mirror

Port 1

ACL ips

Apply Revert

To show the ACLs to be mirrored:

1. Select Configure Interface from the Step list.
2. Select Show Mirror from the Action list.
3. Select a port.

Figure 166: Showing the VLANs to Mirror

Security > ACL

Step: 3. Configure Interface Action: Show Mirror

Port 1

ACL Mirror List Total: 1

<input type="checkbox"/>	Source Access-List
<input type="checkbox"/>	ips

Delete Revert

Showing ACL Hardware Counters Use the Security > ACL > (Configure Interface - Show Hardware Counters) page to show statistics for ACL hardware counters.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier.
- ◆ **Type** – Selects the type of ACL.
- ◆ **Direction** – Displays statistics for ingress or egress traffic.
- ◆ **Time Range** – Name of a time range.
- ◆ **Counter** – Activates the counter for specified ACL.
- ◆ **Query** – Displays statistics for selected criteria.
- ◆ **ACL Name** – The ACL bound to this port.
- ◆ **Action** – Shows if action is to permit or deny specified packets.
- ◆ *Rules* – Shows the rules for the ACL bound to this port.
- ◆ **Time Range** – Name of a time range.
- ◆ **Hit** – Shows the number of packets matching this ACL.⁷
- ◆ **Clear Counter** – Clears the hit counter for the specified ACL.

Web Interface

To show statistics for ACL hardware counters:

1. Click Security, ACL.
2. Select Configure Interface from the Step list.
3. Select Show Hardware Counters from the Action list.
4. Select a port.
5. Select ingress or egress traffic.

7. Due to a hardware limitation, statistics are only displayed for permit rules.

Figure 167: Showing ACL Statistics

The screenshot shows the 'Security > ACL' configuration page. At the top, it indicates 'Step: 3. Configure Interface' and 'Action: Show Hardware Counter'. Below this, there are dropdown menus for 'Port' (set to 1), 'Type' (set to IP Standard), and 'Direction' (set to Ingress). A 'Query' button is located below these settings. The 'ACL Name' is 'rd'. Below the settings, it says 'Total: 1'. A table displays the ACL statistics:

Action	Source IP Address	Time-Range	Hit	Clear Counter
Permit	Any		14	Clear

ARP Inspection

ARP Inspection is a security feature that validates the MAC Address bindings for Address Resolution Protocol packets. It provides protection against ARP traffic with invalid MAC-to-IP address bindings, which forms the basis for certain “man-in-the-middle” attacks. This is accomplished by intercepting all ARP requests and responses and verifying each of these packets before the local ARP cache is updated or the packet is forwarded to the appropriate destination. Invalid ARP packets are dropped.

ARP Inspection determines the validity of an ARP packet based on valid IP-to-MAC address bindings stored in a trusted database – the DHCP snooping binding database (see [“DHCP Snooping Global Configuration” on page 320](#)). This database is built by DHCP snooping if it is enabled on globally on the switch and on the required VLANs. ARP Inspection can also validate ARP packets against user-configured ARP access control lists (ACLs) for hosts with statically configured addresses (see [“Configuring an ARP ACL” on page 279](#)).

Command Usage

Enabling & Disabling ARP Inspection

- ◆ ARP Inspection is controlled on a global and VLAN basis.
- ◆ By default, ARP Inspection is disabled both globally and on all VLANs.
 - If ARP Inspection is globally enabled, then it becomes active only on the VLANs where it has been enabled.
 - When ARP Inspection is enabled globally, all ARP request and reply packets on inspection-enabled VLANs are redirected to the CPU and their switching behavior handled by the ARP Inspection engine.
 - If ARP Inspection is disabled globally, then it becomes inactive for all VLANs, including those where inspection is enabled.

- When ARP Inspection is disabled, all ARP request and reply packets will bypass the ARP Inspection engine and their switching behavior will match that of all other packets.
- Disabling and then re-enabling global ARP Inspection will not affect the ARP Inspection configuration of any VLANs.
- When ARP Inspection is disabled globally, it is still possible to configure ARP Inspection for individual VLANs. These configuration changes will only become active after ARP Inspection is enabled globally again.
- ◆ The ARP Inspection engine in the current firmware version does not support ARP Inspection on trunk ports.

Configuring Global Settings for ARP Inspection

Use the Security > ARP Inspection (Configure General) page to enable ARP inspection globally for the switch, to validate address information in each packet, and configure logging.

Command Usage

ARP Inspection Validation

- ◆ By default, ARP Inspection Validation is disabled.
- ◆ Specifying at least one of the following validations enables ARP Inspection Validation globally. Any combination of the following checks can be active concurrently.
 - Destination MAC – Checks the destination MAC address in the Ethernet header against the target MAC address in the ARP body. This check is performed for ARP responses. When enabled, packets with different MAC addresses are classified as invalid and are dropped.
 - IP – Checks the ARP body for invalid and unexpected IP addresses. These addresses include 0.0.0.0, 255.255.255.255, and all IP multicast addresses. Sender IP addresses are checked in all ARP requests and responses, while target IP addresses are checked only in ARP responses.
 - Source MAC – Checks the source MAC address in the Ethernet header against the sender MAC address in the ARP body. This check is performed on both ARP requests and responses. When enabled, packets with different MAC addresses are classified as invalid and are dropped.

ARP Inspection Logging

- ◆ By default, logging is active for ARP Inspection, and cannot be disabled.
- ◆ The administrator can configure the log facility rate.

- ◆ When the switch drops a packet, it places an entry in the log buffer, then generates a system message on a rate-controlled basis. After the system message is generated, the entry is cleared from the log buffer.
- ◆ Each log entry contains flow information, such as the receiving VLAN, the port number, the source and destination IP addresses, and the source and destination MAC addresses.
- ◆ If multiple, identical invalid ARP packets are received consecutively on the same VLAN, then the logging facility will only generate one entry in the log buffer and one corresponding system message.
- ◆ If the log buffer is full, the oldest entry will be replaced with the newest entry.

Parameters

These parameters are displayed:

- ◆ **ARP Inspection Status** – Enables ARP Inspection globally. (Default: Disabled)
- ◆ **ARP Inspection Validation** – Enables extended ARP Inspection Validation if any of the following options are enabled. (Default: Disabled)
 - **Dst-MAC** – Validates the destination MAC address in the Ethernet header against the target MAC address in the body of ARP responses.
 - **IP** – Checks the ARP body for invalid and unexpected IP addresses. Sender IP addresses are checked in all ARP requests and responses, while target IP addresses are checked only in ARP responses.
 - **Allow Zeros** – Allows sender IP address to be 0.0.0.0.
 - **Src-MAC** – Validates the source MAC address in the Ethernet header against the sender MAC address in the ARP body. This check is performed on both ARP requests and responses.
- ◆ **Log Message Number** – The maximum number of entries saved in a log message. (Range: 0-256; Default: 5)
- ◆ **Log Interval** – The interval at which log messages are sent. (Range: 0-86400 seconds; Default: 1 second)

Web Interface

To configure global settings for ARP Inspection:

1. Click Security, ARP Inspection.
2. Select Configure General from the Step list.
3. Enable ARP inspection globally, enable any of the address validation options, and adjust any of the logging parameters if required.

4. Click Apply.

Figure 168: Configuring Global Settings for ARP Inspection

Security > ARP Inspection

Step: 1. Configure General

ARP Inspection Status Enabled

ARP Inspection Validation Dst-MAC IP Src-MAC

Log Message Number (0-256) 50

Log Interval (0-86400) 100 sec

Apply Revert

Configuring VLAN Settings for ARP Inspection

Use the Security > ARP Inspection (Configure VLAN) page to enable ARP inspection for any VLAN and to specify the ARP ACL to use.

Command Usage

ARP Inspection VLAN Filters (ACLs)

- ◆ By default, no ARP Inspection ACLs are configured and the feature is disabled.
- ◆ ARP Inspection ACLs are configured within the ARP ACL configuration page (see [page 279](#)).
- ◆ ARP Inspection ACLs can be applied to any configured VLAN.
- ◆ ARP Inspection uses the DHCP snooping bindings database for the list of valid IP-to-MAC address bindings. ARP ACLs take precedence over entries in the DHCP snooping bindings database. The switch first compares ARP packets to any specified ARP ACLs.
- ◆ If *Static* is specified, ARP packets are only validated against the selected ACL – packets are filtered according to any matching rules, packets not matching any rules are dropped, and the DHCP snooping bindings database check is bypassed.
- ◆ If *Static* is not specified, ARP packets are first validated against the selected ACL; if no ACL rules match the packets, then the DHCP snooping bindings database determines their validity.

Parameters

These parameters are displayed:

- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)

- ◆ **DAI Status** – Enables Dynamic ARP Inspection for the selected VLAN. (Default: Disabled)
- ◆ **ACL Name** – Allows selection of any configured ARP ACLs. (Default: None)
- ◆ **Static** – When an ARP ACL is selected, and static mode also selected, the switch only performs ARP Inspection and bypasses validation against the DHCP Snooping Bindings database. When an ARP ACL is selected, but static mode is not selected, the switch first performs ARP Inspection and then validation against the DHCP Snooping Bindings database. (Default: Disabled)

Web Interface

To configure VLAN settings for ARP Inspection:

1. Click Security, ARP Inspection.
2. Select Configure VLAN from the Step list.
3. Enable ARP inspection for the required VLANs, select an ARP ACL filter to check for configured addresses, and select the Static option to bypass checking the DHCP snooping bindings database if required.
4. Click Apply.

Figure 169: Configuring VLAN Settings for ARP Inspection

VLAN	DAI Status	ACL Name	ACL Status
1	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> R&D	<input type="checkbox"/> Static
2	<input type="checkbox"/> Enabled	<input type="checkbox"/> R&D	<input type="checkbox"/> Static
3	<input type="checkbox"/> Enabled	<input type="checkbox"/> R&D	<input type="checkbox"/> Static
4	<input type="checkbox"/> Enabled	<input type="checkbox"/> R&D	<input type="checkbox"/> Static
5	<input type="checkbox"/> Enabled	<input type="checkbox"/> R&D	<input type="checkbox"/> Static

Configuring Interface Settings for ARP Inspection

Use the Security > ARP Inspection (Configure Interface) page to specify the ports that require ARP inspection, and to adjust the packet inspection rate.

Parameters

These parameters are displayed:

- ◆ **Interface** – Port or trunk identifier.
- ◆ **Trust Status** – Configures the port as trusted or untrusted. (Default: Untrusted)
By default, all untrusted ports are subject to ARP packet rate limiting, and all trusted ports are exempt from ARP packet rate limiting.

Packets arriving on trusted interfaces bypass all ARP Inspection and ARP Inspection Validation checks and will always be forwarded, while those arriving on untrusted interfaces are subject to all configured ARP inspection tests.

- ◆ **Packet Rate Limit** – Sets the maximum number of ARP packets that can be processed by the CPU per second on trusted or untrusted ports. (Range: 0-2048; Default: 15)

Setting the rate limit to “0” means that there is no restriction on the number of ARP packets that can be processed by the CPU.

The switch will drop all ARP packets received on a port which exceeds the configured ARP-packets-per-second rate limit.

Web Interface

To configure interface settings for ARP Inspection:

1. Click Security, ARP Inspection.
2. Select Configure Interface from the Step list.
3. Specify any untrusted ports which require ARP inspection, and adjust the packet inspection rate.
4. Click Apply.

Figure 170: Configuring Interface Settings for ARP Inspection

Port	Trust Status	Packet Rate Limit (0-2048 pps)
1	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> 15
2	<input type="checkbox"/> Enabled	<input checked="" type="checkbox"/> 15
3	<input type="checkbox"/> Enabled	<input checked="" type="checkbox"/> 15
4	<input type="checkbox"/> Enabled	<input checked="" type="checkbox"/> 15
5	<input type="checkbox"/> Enabled	<input checked="" type="checkbox"/> 15

Displaying ARP Inspection Statistics Use the Security > ARP Inspection (Show Information - Show Statistics) page to display statistics about the number of ARP packets processed, or dropped for various reasons.

Parameters

These parameters are displayed:

Table 23: ARP Inspection Statistics

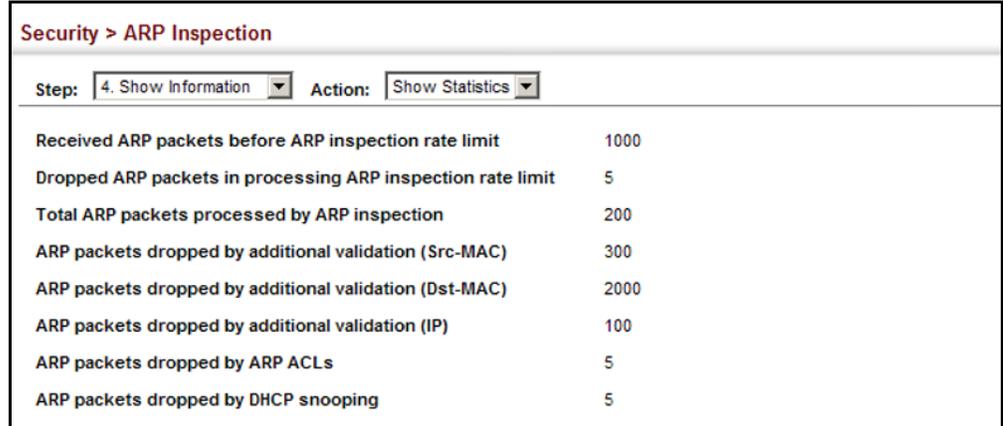
Parameter	Description
Received ARP packets before ARP inspection rate limit	Count of ARP packets received but not exceeding the ARP Inspection rate limit.
Dropped ARP packets in the process of ARP inspection rate limit	Count of ARP packets exceeding (and dropped by) ARP rate limiting.
Total ARP packets processed by ARP inspection	Count of all ARP packets processed by the ARP Inspection engine.
ARP packets dropped by additional validation (Src-MAC)	Count of packets that failed the source MAC address test.
ARP packets dropped by additional validation (Dst-MAC)	Count of packets that failed the destination MAC address test.
ARP packets dropped by additional validation (IP)	Count of ARP packets that failed the IP address test.
ARP packets dropped by ARP ACLs	Count of ARP packets that failed validation against ARP ACL rules.
ARP packets dropped by DHCP snooping	Count of packets that failed validation against the DHCP Snooping Binding database.

Web Interface

To display statistics for ARP Inspection:

1. Click Security, ARP Inspection.
2. Select Show Information from the Step list.
3. Select Show Statistics from the Action list.

Figure 171: Displaying Statistics for ARP Inspection



The screenshot shows the 'Security > ARP Inspection' page. At the top, there is a breadcrumb trail 'Security > ARP Inspection'. Below it, there are two dropdown menus: 'Step: 4. Show Information' and 'Action: Show Statistics'. The main content is a table with two columns: the first column lists various ARP-related statistics, and the second column shows their corresponding counts.

Security > ARP Inspection	
Step:	4. Show Information
Action:	Show Statistics
Received ARP packets before ARP inspection rate limit	1000
Dropped ARP packets in processing ARP inspection rate limit	5
Total ARP packets processed by ARP inspection	200
ARP packets dropped by additional validation (Src-MAC)	300
ARP packets dropped by additional validation (Dst-MAC)	2000
ARP packets dropped by additional validation (IP)	100
ARP packets dropped by ARP ACLs	5
ARP packets dropped by DHCP snooping	5

Displaying the ARP Inspection Log

Use the Security > ARP Inspection (Show Information - Show Log) page to show information about entries stored in the log, including the associated VLAN, port, and address components.

Parameters

These parameters are displayed:

Table 24: ARP Inspection Log

Parameter	Description
VLAN ID	The VLAN where this packet was seen.
Port	The port where this packet was seen.
Src. IP Address	The source IP address in the packet.
Dst. IP Address	The destination IP address in the packet.
Src. MAC Address	The source MAC address in the packet.
Dst. MAC Address	The destination MAC address in the packet.

Web Interface

To display the ARP Inspection log:

1. Click Security, ARP Inspection.
2. Select Show Information from the Step list.
3. Select Show Log from the Action list.

Figure 172: Displaying the ARP Inspection Log

Security > ARP Inspection					
Step: 4. Show Information		Action: Show Log			
ARP Inspection Log List					Total: 2
VLAN ID	Port	Src. IP Address	Dst. IP Address	Src. MAC Address	Dst. MAC Address
1	15	192.168.1.1	192.168.1.5	11-22-33-44-55-66	AA-BB-CC-DD-EE-FF
1	17	192.168.1.3	192.168.1.23	11-4E-33-75-55-BB	A0-3B-C9-DD-4E-1F

Filtering IP Addresses for Management Access

Use the Security > IP Filter page to create a list of up to 15 IP addresses or IP address groups that are allowed management access to the switch through the web interface, SNMP, or Telnet.

Command Usage

- ◆ The management interfaces are open to all IP addresses by default. Once you add an entry to a filter list, access to that interface is restricted to the specified addresses.
- ◆ If anyone tries to access a management interface on the switch from an invalid address, the switch will reject the connection, enter an event message in the system log, and send a trap message to the trap manager.
- ◆ IP address can be configured for SNMP, web and Telnet access respectively. Each of these groups can include up to five different sets of addresses, either individual addresses or address ranges.
- ◆ When entering addresses for the same group (i.e., SNMP, web or Telnet), the switch will not accept overlapping address ranges. When entering addresses for different groups, the switch will accept overlapping address ranges.
- ◆ You cannot delete an individual address from a specified range. You must delete the entire range, and reenter the addresses.
- ◆ You can delete an address range just by specifying the start address, or by specifying both the start address and end address.

Parameters

These parameters are displayed:

- ◆ **Mode**
 - **Web** – Configures IP address(es) for the web group.
 - **SNMP** – Configures IP address(es) for the SNMP group.
 - **Telnet** – Configures IP address(es) for the Telnet group.

- **All** – Configures IP address(es) for all groups.
- ◆ **Start IP Address** – A single IP address, or the starting address of a range.
- ◆ **End IP Address** – The end address of a range.

Web Interface

To create a list of IP addresses authorized for management access:

1. Click Security, IP Filter.
2. Select Add from the Action list.
3. Select the management interface to filter (Web, SNMP, Telnet, All).
4. Enter the IP addresses or range of addresses that are allowed management access to an interface.
5. Click Apply

Figure 173: Creating an IP Address Filter for Management Access

Security > IP Filter

Action: Add

Mode Web SNMP Telnet All

Start IP Address 10.1.2.3

End IP Address

Apply Revert

To show a list of IP addresses authorized for management access:

1. Click Security, IP Filter.
2. Select Show from the Action list.

Figure 174: Showing IP Addresses Authorized for Management Access

Security > IP Filter

Action: Show

Mode Web SNMP Telnet All

SNMP IP Filter List Total: 1

<input type="checkbox"/>	Start IP Address	End IP Address
<input type="checkbox"/>	10.1.2.3	10.1.2.3

Delete Revert

Configuring Port Security

Use the Security > Port Security page to configure the maximum number of device MAC addresses that can be learned by a switch port, stored in the address table, and authorized to access the network.

When port security is enabled on a port, the switch stops learning new MAC addresses on the specified port when it has reached a configured maximum number. Only incoming traffic with source addresses already stored in the address table will be authorized to access the network through that port. If a device with an unauthorized MAC address attempts to use the switch port, the intrusion will be detected and the switch can automatically take action by disabling the port and sending a trap message.

Command Usage

- ◆ The default maximum number of MAC addresses allowed on a secure port is zero (that is, disabled). To use port security, you must configure the maximum number of addresses allowed on a port.
- ◆ To configure the maximum number of address entries which can be learned on a port, and then specify the maximum number of dynamic addresses allowed. The switch will learn up to the maximum number of allowed address pairs <source MAC address, VLAN> for frames received on the port. When the port has reached the maximum number of MAC addresses, the port will stop learning new addresses. The MAC addresses already in the address table will be retained and will not be aged out.

Note that you can manually add additional secure addresses to a port using the Static Address Table ([page 163](#)).

- ◆ When the port security state is changed from enabled to disabled, all dynamically learned entries are cleared from the address table.
- ◆ If port security is enabled, and the maximum number of allowed addresses are set to a non-zero value, any device not in the address table that attempts to use the port will be prevented from accessing the switch.
- ◆ If a port is disabled (shut down) due to a security violation, it must be manually re-enabled from the Interface > Port > General page ([page 103](#)).
- ◆ A secure port has the following restrictions:
 - It cannot be used as a member of a static or dynamic trunk.
 - It should not be connected to a network interconnection device.
 - RSPAN and port security are mutually exclusive functions. If port security is enabled on a port, that port cannot be set as an RSPAN uplink port, ~~source port, or destination port~~. Also, when a port is configured as an RSPAN uplink port, source port, or destination port, port security cannot be enabled on that port.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier.
- ◆ **Security Status** – Enables or disables port security on a port.
(Default: Disabled)
- ◆ **Port Status** – The operational status:
 - Secure/Down – Port security is disabled.
 - Secure/Up – Port security is enabled.
 - Shutdown – Port is shut down due to a response to a port security violation.
- ◆ **Action** – Indicates the action to be taken when a port security violation is detected:
 - **None:** No action should be taken. (This is the default.)
 - **Trap:** Send an SNMP trap message.
 - **Shutdown:** Disable the port.
 - **Trap and Shutdown:** Send an SNMP trap message and disable the port.
- ◆ **Max MAC Count** – The maximum number of MAC addresses that can be learned on a port. (Range: 0-1024, where 0 means disabled)
The maximum address count is effective when port security is enabled or disabled.
- ◆ **Current MAC Count** – The number of MAC addresses currently associated with this interface.
- ◆ **MAC Filter** – Shows if MAC address filtering has been set under Security > Network Access (Configure MAC Filter) as described on [page 252](#).
- ◆ **MAC Filter ID** – The identifier for a MAC address filter.
- ◆ **Last Intrusion MAC** – The last unauthorized MAC address detected.
- ◆ **Last Time Detected Intrusion MAC** – The last time an unauthorized MAC address was detected.

Web Interface

To configure port security:

1. Click Security, Port Security.
2. Mark the check box in the Security Status column to enable security, set the action to take when an invalid address is detected on a port, and set the maximum number of MAC addresses allowed on the port.
3. Click Apply.

Figure 175: Configuring Port Security

Port	Security Status	Port Status	Action	Max MAC Count (0-1024)	Current MAC Count	MAC Filter	MAC Filter ID	Last Intrusion MAC	Last Time Detected Intrusion MAC
1	<input checked="" type="checkbox"/> Enabled	Secure/Down	Trap and Shutdown	0	0	Disabled	0	NA	NA
2	<input type="checkbox"/> Enabled	Secure/Down	None	0	0	Disabled	0	NA	NA
3	<input type="checkbox"/> Enabled	Secure/Down	None	0	0	Disabled	0	NA	NA
4	<input type="checkbox"/> Enabled	Secure/Down	None	0	0	Disabled	0	NA	NA
5	<input type="checkbox"/> Enabled	Secure/Down	None	0	0	Disabled	0	NA	NA

Configuring 802.1X Port Authentication

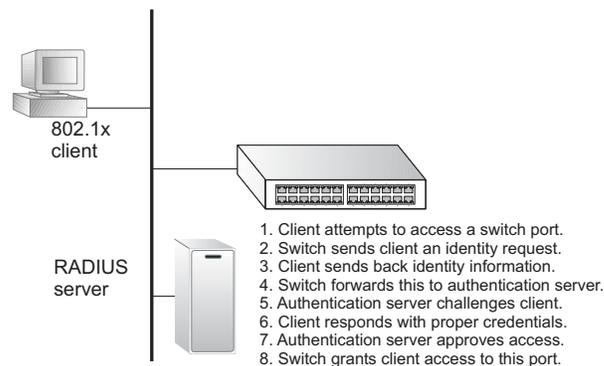
Network switches can provide open and easy access to network resources by simply attaching a client PC. Although this automatic configuration and access is a desirable feature, it also allows unauthorized personnel to easily intrude and possibly gain access to sensitive network data.

The IEEE 802.1X (dot1X) standard defines a port-based access control procedure that prevents unauthorized access to a network by requiring users to first submit credentials for authentication. Access to all switch ports in a network can be centrally controlled from a server, which means that authorized users can use the same credentials for authentication from any point within the network.

This switch uses the Extensible Authentication Protocol over LANs (EAPOL) to exchange authentication protocol messages with the client, and a remote RADIUS authentication server to verify user identity and access rights. When a client (i.e., Supplicant) connects to a switch port, the switch (i.e., Authenticator) responds with an EAPOL identity request. The client provides its identity (such as a user name) in an EAPOL response to the switch, which it forwards to the RADIUS server. The RADIUS server verifies the client identity and sends an access challenge back to the client. The EAP packet from the RADIUS server contains not only the challenge, but the authentication method to be used. The client can reject the authentication method and request another, depending on the configuration of the client software and the RADIUS server. The encryption method used to pass authentication messages can be MD5 (Message-Digest 5), TLS (Transport Layer

Security), PEAP (Protected Extensible Authentication Protocol), or TTLS (Tunneled Transport Layer Security). The client responds to the appropriate method with its credentials, such as a password or certificate. The RADIUS server verifies the client credentials and responds with an accept or reject packet. If authentication is successful, the switch allows the client to access the network. Otherwise, non-EAP traffic on the port is blocked or assigned to a guest VLAN based on the “intrusion-action” setting. In “multi-host” mode, only one host connected to a port needs to pass authentication for all other hosts to be granted network access. Similarly, a port can become unauthorized for all hosts if one attached host fails re-authentication or sends an EAPOL logoff message.

Figure 176: Configuring Port Security



The operation of 802.1X on the switch requires the following:

- ◆ The switch must have an IP address assigned.
- ◆ RADIUS authentication must be enabled on the switch and the IP address of the RADIUS server specified.
- ◆ 802.1X must be enabled globally for the switch.
- ◆ Each switch port that will be used must be set to dot1X “Auto” mode.
- ◆ Each client that needs to be authenticated must have dot1X client software installed and properly configured.
- ◆ The RADIUS server and 802.1X client support EAP. (The switch only supports EAPOL in order to pass the EAP packets from the server to the client.)
- ◆ The RADIUS server and client also have to support the same EAP authentication type – MD5, PEAP, TLS, or TTLS. Native support for these encryption methods is provided in Windows 7, 8 and 10.

Configuring 802.1X Global Settings

Use the Security > Port Authentication (Configure Global) page to configure IEEE 802.1X port authentication. The 802.1X protocol must be enabled globally for the switch system before port settings are active.

Parameters

These parameters are displayed:

- ◆ **System Authentication Control** – Sets the global setting for 802.1X. (Default: Disabled)

- ◆ **EAPOL Pass Through** – Passes EAPOL frames through to all ports in STP forwarding state when dot1x is globally disabled. (Default: Disabled)

When this device is functioning as intermediate node in the network and does not need to perform dot1x authentication, **EAPOL Pass Through** can be enabled to allow the switch to forward EAPOL frames from other switches on to the authentication servers, thereby allowing the authentication process to still be carried out by switches located on the edge of the network.

When this device is functioning as an edge switch but does not require any attached clients to be authenticated, **EAPOL Pass Through** can be disabled to discard unnecessary EAPOL traffic.

- ◆ **Default** – Sets all configurable 802.1X global and port settings to their default values.

Web Interface

To configure global settings for 802.1X:

1. Click Security, Port Authentication.
2. Select Configure Global from the Step list.
3. Enable 802.1X globally for the switch, and configure EAPOL Pass Through if required.
4. Click Apply

Figure 177: Configuring Global Settings for 802.1X Port Authentication

The screenshot shows the 'Security > Port Authentication' configuration page. At the top, the breadcrumb 'Security > Port Authentication' is visible. Below it, a 'Step:' dropdown menu is set to '1. Configure Global'. The main configuration area contains two settings: 'System Authentication Control' and 'EAPOL Pass Through', both with 'Enabled' checkboxes that are currently unchecked. At the bottom right of the configuration area are 'Apply' and 'Revert' buttons. At the bottom left, there is a 'Default' button with a tooltip that reads 'Click this button to set 802.1X global/port settings to default values.'

Configuring Port Authenticator Settings for 802.1X

Use the Security > Port Authentication (Configure Interface) page to configure 802.1X port settings for the switch as the local authenticator. When 802.1X is enabled, you need to configure the parameters for the authentication process that runs between the client and the switch (i.e., authenticator), as well as the client identity lookup process that runs between the switch and authentication server.

Command Usage

- ◆ When the switch functions as a local authenticator between supplicant devices attached to the switch and the authentication server, configure the parameters for the exchange of EAP messages between the authenticator and clients on the Authenticator configuration page.
- ◆ This switch can be configured to serve as the authenticator on selected ports by setting the Control Mode to Auto on this configuration page.

Parameters

These parameters are displayed:

- ◆ **Port** – Port number.
- ◆ **Status** – Indicates if authentication is enabled or disabled on the port. The status is disabled if the control mode is set to Force-Authorized.
- ◆ **Authorized** – Displays the 802.1X authorization status of connected clients.
 - **Yes** – Connected client is authorized.
 - **N/A** – Connected client is not authorized, or port is not connected.
- ◆ **Control Mode** – Sets the authentication mode to one of the following options:
 - **Auto** – Requires a dot1x-aware client to be authorized by the authentication server. Clients that are not dot1x-aware will be denied access.
 - **Force-Authorized** – Forces the port to grant access to all clients, either dot1x-aware or otherwise. (This is the default setting.)
 - **Force-Unauthorized** – Forces the port to deny access to all clients, either dot1x-aware or otherwise.
- ◆ **Operation Mode** – Allows single or multiple hosts (clients) to connect to an 802.1X-authorized port. (Default: Single-Host)
 - **Single-Host** – Allows only a single host to connect to this port.
 - **Multi-Host** – Allows multiple host to connect to this port.
In this mode, only one host connected to a port needs to pass authentication for all other hosts to be granted network access. Similarly, a

port can become unauthorized for all hosts if one attached host fails re-authentication or sends an EAPOL logoff message.

- **MAC-Based** – Allows multiple hosts to connect to this port, with each host needing to be authenticated.

In this mode, each host connected to a port needs to pass authentication. The number of hosts allowed access to a port operating in this mode is limited only by the available space in the secure address table (i.e., up to 1024 addresses).

- ◆ **Max Count** – The maximum number of hosts that can connect to a port when the Multi-Host operation mode is selected. (Range: 1-1024; Default: 5)
- ◆ **Max Request** – Sets the maximum number of times the switch port will retransmit an EAP request packet to the client before it times out the authentication session. (Range: 1-10; Default 2)
- ◆ **Quiet Period** – Sets the time that a switch port waits after the Max Request Count has been exceeded before attempting to acquire a new client. (Range: 1-65535 seconds; Default: 60 seconds)
- ◆ **Tx Period** – Sets the time period during an authentication session that the switch waits before re-transmitting an EAP packet. (Range: 1-65535; Default: 30 seconds)
- ◆ **Supplicant Timeout** – Sets the time that a switch port waits for a response to an EAP request from a client before re-transmitting an EAP packet. (Range: 1-65535; Default: 30 seconds)

This command attribute sets the timeout for EAP-request frames other than EAP-request/identity frames. If dot1x authentication is enabled on a port, the switch will initiate authentication when the port link state comes up. It will send an EAP-request/identity frame to the client to request its identity, followed by one or more requests for authentication information. It may also send other EAP-request frames to the client during an active connection as required for reauthentication.

- ◆ **Server Timeout** – Sets the time that a switch port waits for a response to an EAP request from an authentication server before re-transmitting an EAP packet. (Default: 0 seconds)

A RADIUS server must be set before the correct operational value of 10 seconds will be displayed in this field. (See [“Configuring Remote Logon Authentication Servers” on page 238.](#))
- ◆ **Re-authentication Status** – Sets the client to be re-authenticated after the interval specified by the Re-authentication Period. Re-authentication can be used to detect if a new device is plugged into a switch port. (Default: Disabled)

- ◆ **Re-authentication Period** – Sets the time period after which a connected client must be re-authenticated. (Range: 1-65535 seconds; Default: 3600 seconds)
- ◆ **Re-authentication Max Retries** – The maximum number of times the switch port will retransmit an EAP request/identity packet to the client before it times out the authentication session. (Range: 1-10; Default: 2)
- ◆ **Intrusion Action** – Sets the port's response to a failed authentication.
 - **Block Traffic** – Blocks all non-EAP traffic on the port. (This is the default setting.)
 - **Guest VLAN** – All traffic for the port is assigned to a guest VLAN. The guest VLAN must be separately configured (See [“Configuring VLAN Groups” on page 147](#)) and mapped on each port (See [“Configuring Network Access for Ports” on page 249](#)).

Supplicant List

- ◆ **Supplicant** – MAC address of authorized client.

Authenticator PAE State Machine

- ◆ **State** – Current state (including initialize, disconnected, connecting, authenticating, authenticated, aborting, held, force_authorized, force_unauthorized).
- ◆ **Reauth Count** – Number of times connecting state is re-entered.
- ◆ **Current Identifier** – Identifier sent in each EAP Success, Failure or Request packet by the Authentication Server.

Backend State Machine

- ◆ **State** – Current state (including request, response, success, fail, timeout, idle, initialize).
- ◆ **Request Count** – Number of EAP Request packets sent to the Supplicant without receiving a response.
- ◆ **Identifier (Server)** – Identifier carried in the most recent EAP Success, Failure or Request packet received from the Authentication Server.

Reauthentication State Machine

- ◆ **State** – Current state (including initialize, reauthenticate).

Web Interface

To configure port authenticator settings for 802.1X:

1. Click Security, Port Authentication.
2. Select Configure Interface from the Step list.
3. Modify the authentication settings for each port as required.
4. Click Apply

Figure 178: Configuring Interface Settings for 802.1X Port Authenticator

Security > Port Authentication

Step: 2. Configure Interface

Port: 1

Status: Disabled

Authorized: Yes

Control Mode: Auto

Operation Mode: Single-Host

Max MAC Count (1-1024): 5

Max Request (1-10): 2

Quiet Period (1-65535): 60 sec

Tx Period (1-65535): 30 sec

Supplicant Timeout (1-65535): 30 sec

Server Timeout: 10 sec

Re-authentication Status: Enabled

Re-authentication Period (1-65535): 3600 sec

Re-authentication Max Retries (1-10): 2

Intrusion Action: Block Traffic

Supplicant List Total: 1

Supplicant	Authenticator PAE State Machine			Backend State Machine			Reauthentication State Machine
	State	Reauth Count	Current Identifier	State	Request Count	Identifier (Server)	State
00-00-00-00-00-00	Initialize	0	0	Initialize	0	0	Initialize

Apply
Revert

Displaying 802.1X Statistics Use the Security > Port Authentication (Show Statistics) page to display statistics for dot1x protocol exchanges for any port.

Parameters

These parameters are displayed:

Table 25: 802.1X Statistics

Parameter	Description
<i>Authenticator</i>	
Rx EAPOL Start	The number of EAPOL Start frames that have been received by this Authenticator.
Rx EAPOL Logoff	The number of EAPOL Logoff frames that have been received by this Authenticator.
Rx EAPOL Invalid	The number of EAPOL frames that have been received by this Authenticator in which the frame type is not recognized.
Rx EAPOL Total	The number of valid EAPOL frames of any type that have been received by this Authenticator.
Rx Last EAPOLVer	The protocol version number carried in the most recent EAPOL frame received by this Authenticator.
Rx Last EAPOLSrc	The source MAC address carried in the most recent EAPOL frame received by this Authenticator.
Rx EAP Resp/Id	The number of EAP Resp/Id frames that have been received by this Authenticator.
Rx EAP Resp/Oth	The number of valid EAP Response frames (other than Resp/Id frames) that have been received by this Authenticator.
Rx EAP LenError	The number of EAPOL frames that have been received by this Authenticator in which the Packet Body Length field is invalid.
Tx EAP Req/Id	The number of EAP Req/Id frames that have been transmitted by this Authenticator.
Tx EAP Req/Oth	The number of EAP Request frames (other than Rq/Id frames) that have been transmitted by this Authenticator.
Tx EAPOL Total	The number of EAPOL frames of any type that have been transmitted by this Authenticator.
<i>Supplicant</i>	
Rx EAPOL Invalid	The number of EAPOL frames that have been received by this Supplicant in which the frame type is not recognized.
Rx EAPOL Total	The number of valid EAPOL frames of any type that have been received by this Supplicant.
Rx Last EAPOLVer	The protocol version number carried in the most recent EAPOL frame received by this Supplicant.
Rx Last EAPOLSrc	The source MAC address carried in the most recent EAPOL frame received by this Supplicant.
Rx EAP Resp/Id	The number of EAP Resp/Id frames that have been received by this Supplicant.
Rx EAP Resp/Oth	The number of valid EAP Response frames (other than Resp/Id frames) that have been received by this Supplicant.

Table 25: 802.1X Statistics (Continued)

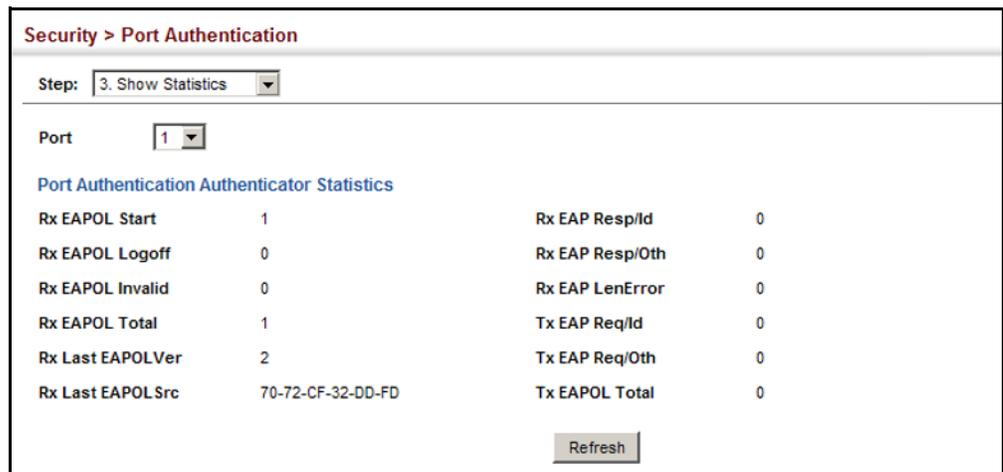
Parameter	Description
Rx EAP LenError	The number of EAPOL frames that have been received by this Supplicant in which the Packet Body Length field is invalid.
Tx EAPOL Total	The number of EAPOL frames of any type that have been transmitted by this Supplicant.
Tx EAPOL Start	The number of EAPOL Start frames that have been transmitted by this Supplicant.
Tx EAPOL Logoff	The number of EAPOL Logoff frames that have been transmitted by this Supplicant.
Tx EAP Req/Id	The number of EAP Req/Id frames that have been transmitted by this Supplicant.
Tx EAP Req/Oth	The number of EAP Request frames (other than Req/Id frames) that have been transmitted by this Supplicant.

Web Interface

To display port authenticator statistics for 802.1X:

1. Click Security, Port Authentication.
2. Select Show Statistics from the Step list.
3. Select a port.

Figure 179: Showing Statistics for 802.1X Port Authenticator



IPv4 Source Guard

IPv4 Source Guard is a security feature that filters IP traffic on network interfaces based on manually configured entries in the IP Source Guard table, or dynamic entries in the DHCP Snooping table when enabled (see [“DHCP Snooping” on page 318](#)). IPv4 source guard can be used to prevent traffic attacks caused when a host tries to use the IPv4 address of a neighbor to access the network. This section describes how to configure IP Source Guard.

Configuring Ports for IPv4 Source Guard

Use the Security > IP Source Guard > Port Configuration page to set the filtering type based on source IP address, or source IP address and MAC address pairs. It also specifies lookup within the ACL binding table or the MAC address binding table, as well as the maximum number of allowed binding entries for the lookup tables.

IP Source Guard is used to filter traffic on an insecure port which receives messages from outside the network or fire wall, and therefore may be subject to traffic attacks caused by a host trying to use the IP address of a neighbor.

Command Usage

- ◆ Setting source guard mode to SIP (Source IP) or SIP-MAC (Source IP and MAC) enables this function on the selected port. Use the SIP option to check the VLAN ID, source IP address, and port number against all entries in the binding table. Use the SIP-MAC option to check these same parameters, plus the source MAC address. If no matching entry is found, the packet is dropped.



Note: Multicast addresses cannot be used by IP Source Guard.

- ◆ When enabled, traffic is filtered based upon dynamic entries learned via DHCP snooping (see [“DHCP Snooping” on page 318](#)), or static addresses configured in the source guard binding table.
- ◆ If IP source guard is enabled, an inbound packet’s IP address (SIP option) or both its IP address and corresponding MAC address (SIP-MAC option) will be checked against the binding table. If no matching entry is found, the packet will be dropped.
- ◆ An entry with same MAC address and a different VLAN ID cannot be added to the binding table .
- ◆ Filtering rules are implemented as follows:
 - If DHCP snooping is disabled (see [page 320](#)), IPv4 source guard will check the VLAN ID, source IP address, port number, and source MAC address (for the SIP-MAC option). If a matching entry is found in the binding table and the entry type is static IPv4 source guard binding, the packet will be forwarded.

- If DHCP snooping is enabled, IPv4 source guard will check the VLAN ID, source IP address, port number, and source MAC address (for the SIP-MAC option). If a matching entry is found in the binding table and the entry type is static IPv4 source guard binding, or dynamic DHCP snooping binding, the packet will be forwarded.
- If IPv4 source guard is enabled on an interface for which IPv4 source bindings have not yet been configured (neither by static configuration in the IPv4 source guard binding table nor dynamically learned from DHCP snooping), the switch will drop all IP traffic on that port, except for DHCP packets allowed by DHCP snooping.

Parameters

These parameters are displayed:

- ◆ **Filter Type** – Configures the switch to filter inbound traffic based source IP address, or source IP address and corresponding MAC address. (Default: None)
 - **Disabled** – Disables IPv4 source guard filtering on the port.
 - **SIP** – Enables traffic filtering based on IP addresses stored in the binding table.
 - **SIP-MAC** – Enables traffic filtering based on IP addresses and corresponding MAC addresses stored in the binding table.
- ◆ **Filter Table** – Sets the source guard learning model to search for addresses in the ACL binding table or the MAC address binding table. (Default: ACL binding table)
- ◆ **Max Binding Entry** – The maximum number of entries that can be bound to an interface. (ACL Table: 1-16, default: 5; MAC Table: 1-1024, default: 1024)

This parameter sets the maximum number of address entries that can be mapped to an interface in the binding table, including both dynamic entries discovered by DHCP snooping (see [“DHCP Snooping” on page 318](#)) and static entries set by IP source guard (see [“Configuring Static Bindings for IP Source Guard” on page 308](#)).

The maximum binding for ACL mode restricts the number of “active” entries per port. If binding entries exceed the maximum number in IPv4 source guard, only the maximum number of binding entries will be set. Dynamic binding entries exceeding the maximum number will be created but will not be active.

The maximum binding for MAC mode restricts the number of MAC addresses learned per port. Authenticated IP traffic with different source MAC addresses cannot be learned if it would exceed this maximum number.

Web Interface

To set the IP Source Guard filter for ports:

1. Click Security, IP Source Guard, Port Configuration.
2. Set the required filtering type, set the table type to use ACL or MAC address binding, and then set the maximum binding entries for each port.
3. Click Apply

Figure 180: Setting the Filter Type for IPv4 Source Guard

Port	Filter Type	Filter Table	ACL Table Max Binding Entry (1-5)	MAC Table Max Binding Entry (1-1024)
1	DISABLED	ACL	5	1024
2	DISABLED	ACL	5	1024
3	DISABLED	ACL	5	1024
4	DISABLED	ACL	5	1024
5	SIP	ACL	5	1024

Configuring Static Bindings for IP Source Guard

Use the Security > IP Source Guard > Static Binding (Configure ACL Table and Configure MAC Table) pages to bind a static address to a port. Table entries include a MAC address, IP address, lease time, entry type (Static, Dynamic), VLAN identifier, and port identifier. All static entries are configured with an infinite lease time, which is indicated with a value of zero in the table.

Command Usage

- ◆ Table entries include a MAC address, IP address, lease time, entry type (Static-IP-SG-Binding, Dynamic-DHCP-Binding), VLAN identifier, and port identifier.
- ◆ Static addresses entered in the source guard binding table are automatically configured with an infinite lease time.
- ◆ When source guard is enabled, traffic is filtered based upon dynamic entries learned via DHCP snooping, or static addresses configured in the source guard binding table.
- ◆ An entry with same MAC address and a different VLAN ID cannot be added to the binding table .
- ◆ Static bindings are processed as follows:
 - A valid static IP source guard entry will be added to the binding table in ACL mode if one of the following conditions is true:
 - If there is no entry with the same VLAN ID and MAC address, a new entry is added to the binding table using the type "static IP source guard binding."

- If there is an entry with the same VLAN ID and MAC address, and the type of entry is static IP source guard binding, then the new entry will replace the old one.
- If there is an entry with the same VLAN ID and MAC address, and the type of the entry is dynamic DHCP snooping binding, then the new entry will replace the old one and the entry type will be changed to static IP source guard binding.
- A valid static IP source guard entry will be added to the binding table in MAC mode if one of the following conditions are true:
 - If there is no binding entry with the same IP address and MAC address, a new entry will be added to the binding table using the type of static IP source guard binding entry.
 - If there is a binding entry with same IP address and MAC address, then the new entry shall replace the old one.
- Only unicast addresses are accepted for static bindings.

Parameters

These parameters are displayed:

Configure ACL Table – Add

- ◆ **Port** – The port to which a static entry is bound.
- ◆ **VLAN** – ID of a configured VLAN (Range: 1-4094)
- ◆ **MAC Address** – A valid unicast MAC address.
- ◆ **IP Address** – A valid unicast IP address, including classful types A, B or C.

Configure MAC Table – Add

- ◆ **MAC Address** – A valid unicast MAC address.
- ◆ **VLAN** – ID of a configured VLAN or a range of VLANs. (Range: 1-4094)
- ◆ **IP Address** – A valid unicast IP address, including classful types A, B or C.
- ◆ **Port** – The port to which a static entry is bound. Specify a physical port number or list of port numbers. Separate nonconsecutive port numbers with a comma and no spaces; or use a hyphen to designate a range of port numbers. (Range: 1-54)

Show

- ◆ **MAC Address** – Physical address associated with the entry.
- ◆ **IP Address** – IP address corresponding to the client.

- ◆ **Lease Time** – The time for which this IP address is leased to the client. (This value is zero for all static addresses.)
- ◆ **VLAN** – VLAN to which this entry is bound.
- ◆ **Interface** – The port to which this entry is bound.

Web Interface

To configure static bindings for IP Source Guard:

1. Click Security, IP Source Guard, Static Binding.
2. Select Configure ACL Table or Configure MAC Table from the Step list.
3. Select Add from the Action list.
4. Enter the required bindings for each port.
5. Click Apply

Figure 181: Configuring Static Bindings for IPv4 Source Guard

Security > IP Source Guard > Static Binding

Step: 1. Configure ACL Table Action: Add

Port 1

VLAN 1

MAC Address 00-10-b5-f4-d0-01

IP Address 10.2.44.96

Apply Revert

To display static bindings for IP Source Guard:

1. Click Security, IP Source Guard, Static Binding.
2. Select Show from the Action list.

Figure 182: Displaying Static Bindings for IP Source Guard

Security > IP Source Guard > Static Binding

Step: 1. Configure ACL Table Action: Show

Static Binding List Total: 1

	MAC Address	IP Address	Lease Time (sec)	VLAN	Interface
<input type="checkbox"/>	00-10-B5-F4-D0-01	10.2.44.96	0	1	Unit 1 / Port 1

Delete Revert

Displaying Information for Dynamic IPv4 Source Guard Bindings

Use the Security > IP Source Guard > Dynamic Binding page to display the source-guard binding table for a selected interface.

Parameters

These parameters are displayed:

Query by

- ◆ **Port** – A port on this switch.
- ◆ **VLAN** – ID of a configured VLAN (Range: 1-4094)
- ◆ **MAC Address** – A valid unicast MAC address.
- ◆ **IP Address** – A valid unicast IP address, including classful types A, B or C.

Dynamic Binding List

- ◆ **VLAN** – VLAN to which this entry is bound.
- ◆ **MAC Address** – Physical address associated with the entry.
- ◆ **Interface** – Port to which this entry is bound.
- ◆ **IP Address** – IP address corresponding to the client.
- ◆ **Lease Time** – The time for which this IP address is leased to the client.
- ◆ **Type** – DHCP-Snooping, BOOTP-Snooping

Web Interface

To display the binding table for IP Source Guard:

1. Click Security, IP Source Guard, Dynamic Binding.
2. Mark the search criteria, and enter the required values.
3. Click Query

Figure 183: Showing the IPv4 Source Guard Binding Table

Security > IP Source Guard > Dynamic Binding

Query by:

- Port: 1
- VLAN: 1
- MAC Address:
- IP Address:

Query

Dynamic Binding List Total: 3

VLAN	MAC Address	Interface	IP Address	Lease Time (sec)
1	00-10-B5-F4-00-01	Unit 1 / Port 2	10.2.44.96	5
1	00-10-B5-F4-00-02	Unit 1 / Port 4	10.2.44.97	25
2	00-10-B5-F4-00-03	Unit 1 / Port 7	10.2.44.98	47

IPv6 Source Guard

IPv6 Source Guard is a security feature that filters IPv6 traffic on non-routed, Layer 2 network interfaces based on manually configured entries in the IPv6 Source Guard table, or dynamic entries in the Neighbor Discovery Snooping table or DHCPv6 Snooping table when either snooping protocol is enabled (refer to the DHCPv6 Snooping commands in the *CLI Reference Guide*). IPv6 source guard can be used to prevent traffic attacks caused when a host tries to use the IPv6 address of a neighbor to access the network. This section describes how to configure IPv6 Source Guard.

Configuring Ports for IPv6 Source Guard Use the Security > IPv6 Source Guard > Port Configuration page to filter inbound traffic based on the source IPv6 address stored in the binding table.

IPv6 Source Guard is used to filter traffic on an insecure port which receives messages from outside the network or fire wall, and therefore may be subject to traffic attacks caused by a host trying to use the IPv6 address of a neighbor.

Command Usage

- ◆ Setting source guard mode to SIP (Source IP) enables this function on the selected port. Use the SIP option to check the VLAN ID, IPv6 global unicast source IP address, and port number against all entries in the binding table.

- ◆ After IPv6 source guard is enabled on an interface, the switch initially blocks all IPv6 traffic received on that interface, except for ND packets allowed by ND snooping and DHCPv6 packets allowed by DHCPv6 snooping. A port access control list (ACL) is applied to the interface. Traffic is then filtered based upon dynamic entries learned via ND snooping or DHCPv6 snooping, or static addresses configured in the source guard binding table. The port allows only IPv6 traffic with a matching entry in the binding table and denies all other IPv6 traffic.
- ◆ Table entries include a MAC address, IPv6 global unicast address, entry type (Static-IPv6-SG-Binding, Dynamic-ND-Binding, Dynamic-DHCPv6-Binding), VLAN identifier, and port identifier.
- ◆ Static addresses entered in the source guard binding table (using the Static Binding page) are automatically configured with an infinite lease time. Dynamic entries learned via DHCPv6 snooping are configured by the DHCPv6 server itself.
- ◆ If IPv6 source guard is enabled, an inbound packet's source IPv6 address will be checked against the binding table. If no matching entry is found, the packet will be dropped.
- ◆ Filtering rules are implemented as follows:
 - If ND snooping and DHCPv6 snooping are disabled, IPv6 source guard will check the VLAN ID, source IPv6 address, and port number. If a matching entry is found in the binding table and the entry type is static IPv6 source guard binding, the packet will be forwarded.
 - If ND snooping or DHCP snooping is enabled, IPv6 source guard will check the VLAN ID, source IP address, and port number. If a matching entry is found in the binding table and the entry type is static IPv6 source guard binding, dynamic ND snooping binding, or dynamic DHCPv6 snooping binding, the packet will be forwarded.
 - If IPv6 source guard is enabled on an interface for which IPv6 source bindings (dynamically learned via ND snooping or DHCPv6 snooping, or manually configured) are not yet configured, the switch will drop all IPv6 traffic on that port, except for ND packets and DHCPv6 packets allowed by DHCPv6 snooping.
 - Only IPv6 global unicast addresses are accepted for static bindings.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier (Range: 1-54)
- ◆ **Filter Type** – Configures the switch to filter inbound traffic based on the following options. (Default: Disabled)
 - **Disabled** – Disables IPv6 source guard filtering on the port.

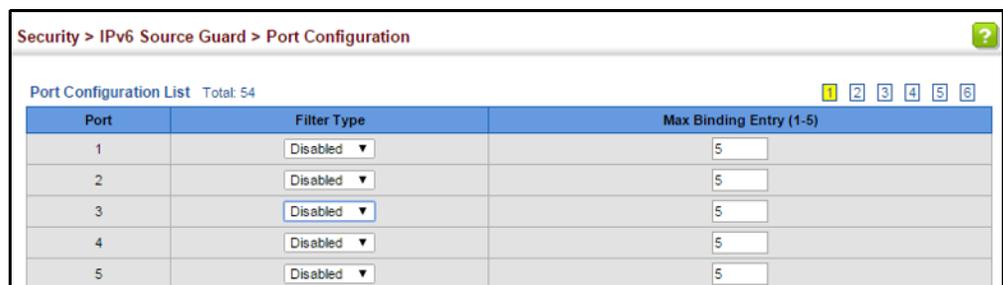
- **SIP** – Enables traffic filtering based on IPv6 global unicast source IPv6 addresses stored in the binding table.
- ◆ **Max Binding Entry** – The maximum number of entries that can be bound to an interface. (Range: 1-5; Default: 5)
 - This parameter sets the maximum number of IPv6 global unicast source IPv6 address entries that can be mapped to an interface in the binding table, including both dynamic entries discovered by ND snooping, DHCPv6 snooping (refer to the DHCPv6 Snooping commands in the *CLI Reference Guide*), and static entries set by IPv6 Source Guard (see “[Configuring Static Bindings for IPv6 Source Guard](#)” on page 315).
 - IPv6 source guard maximum bindings must be set to a value higher than DHCPv6 snooping maximum bindings and ND snooping maximum bindings.
 - If IPv6 source guard, ND snooping, and DHCPv6 snooping are enabled on a port, the dynamic bindings used by ND snooping, DHCPv6 snooping, and IPv6 source guard static bindings cannot exceed the maximum allowed bindings set by this parameter. In other words, no new entries will be added to the IPv6 source guard binding table.
 - If IPv6 source guard is enabled on a port, and the maximum number of allowed bindings is changed to a lower value, precedence is given to deleting entries learned through DHCPv6 snooping, ND snooping, and then manually configured IPv6 source guard static bindings, until the number of entries in the binding table reaches the newly configured maximum number of allowed bindings.

Web Interface

To set the IPv6 Source Guard filter for ports:

1. Click Security, IPv6 Source Guard, Port Configuration.
2. Set the required filtering type for each port.
3. Click Apply

Figure 184: Setting the Filter Type for IPv6 Source Guard



The screenshot shows a web interface for configuring IPv6 Source Guard. The breadcrumb path is "Security > IPv6 Source Guard > Port Configuration". Below the breadcrumb is a "Port Configuration List" with a total of 54 entries. There are six numbered tabs (1-6) and a help icon. The table has three columns: "Port", "Filter Type", and "Max Binding Entry (1-5)".

Port	Filter Type	Max Binding Entry (1-5)
1	Disabled	5
2	Disabled	5
3	Disabled	5
4	Disabled	5
5	Disabled	5

Configuring Static Bindings for IPv6 Source Guard

Use the Security > IPv6 Source Guard > Static Configuration page to bind a static address to a port. Table entries include a MAC address, IPv6 global unicast address, entry type (Static-IPv6-SG-Binding, Dynamic-ND-Binding, Dynamic-DHCPv6-Binding), VLAN identifier, and port identifier.

Command Usage

- ◆ Traffic filtering is based only on the source IPv6 address, VLAN ID, and port number.
- ◆ Static addresses entered in the source guard binding table are automatically configured with an infinite lease time.
- ◆ When source guard is enabled, traffic is filtered based upon dynamic entries learned via ND snooping, DHCPv6 snooping, or static addresses configured in the source guard binding table.
- ◆ An entry with same MAC address and a different VLAN ID cannot be added to the binding table .
- ◆ Static bindings are processed as follows:
 - If there is no entry with same MAC address and IPv6 address, a new entry is added to binding table using static IPv6 source guard binding.
 - If there is an entry with same MAC address and IPv6 address, and the type of entry is static IPv6 source guard binding, then the new entry will replace the old one.
 - If there is an entry with same MAC address and IPv6 address, and the type of the entry is either a dynamic ND snooping binding or DHCPv6 snooping binding, then the new entry will replace the old one and the entry type will be changed to static IPv6 source guard binding.
 - Only unicast addresses are accepted for static bindings.

Parameters

These parameters are displayed:

Add

- ◆ **Port** – The port to which a static entry is bound.
- ◆ **VLAN** – ID of a configured VLAN (Range: 1-4094)
- ◆ **MAC Address** – A valid unicast MAC address.
- ◆ **IPv6 Address** – A valid global unicast IPv6 address. This address must be entered according to RFC 2373 “IPv6 Addressing Architecture,” using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields.

Show

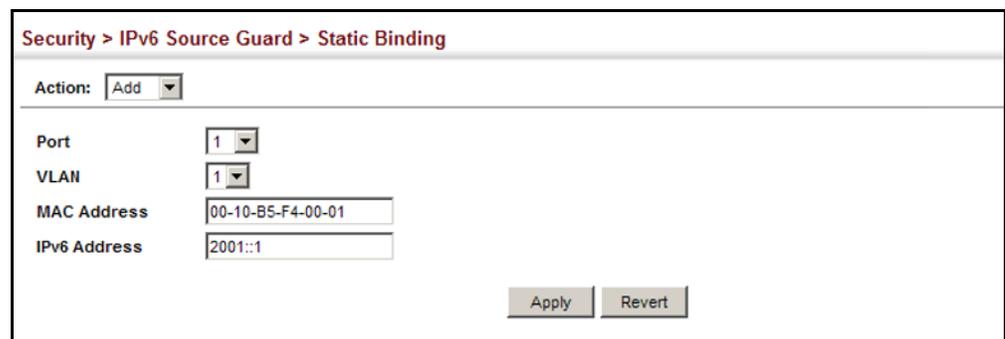
- ◆ **VLAN** – VLAN to which this entry is bound.
- ◆ **MAC Address** – Physical address associated with the entry.
- ◆ **Interface** – The port to which this entry is bound.
- ◆ **IPv6 Address** – IPv6 address corresponding to the client.
- ◆ **Type** – Shows the entry type:
 - **DHCP** – Dynamic DHCPv6 binding, stateful address.
 - **ND** – Dynamic Neighbor Discovery binding, stateless address.
 - **STA** – Static IPv6 Source Guard binding.

Web Interface

To configure static bindings for IPv6 Source Guard:

1. Click Security, IPv6 Source Guard, Static Configuration.
2. Select Add from the Action list.
3. Enter the required bindings for each port.
4. Click Apply

Figure 185: Configuring Static Bindings for IPv6 Source Guard



Security > IPv6 Source Guard > Static Binding

Action: Add

Port: 1

VLAN: 1

MAC Address: 00-10-B5-F4-00-01

IPv6 Address: 2001::1

Apply Revert

To display static bindings for IPv6 Source Guard:

1. Click Security, IPv6 Source Guard, Static Configuration.
2. Select Show from the Action list.

Figure 186: Displaying Static Bindings for IPv6 Source Guard

<input type="checkbox"/>	VLAN	MAC Address	Interface	IPv6 Address	Type
<input type="checkbox"/>	1	00-10-B5-F4-00-01	Eth 1/2	2001:DB8:2222:7272::26	STA
<input type="checkbox"/>	1	00-10-B5-F4-00-02	Eth 1/4	2001:DB8:2222:7272::56	DHCP
<input type="checkbox"/>	2	00-10-B5-F4-00-03	Eth 1/7	2001:DB8:2222:7272::36	ND

Displaying Information for Dynamic IPv6 Source Guard Bindings

Use the Security > IPv6 Source Guard > Dynamic Binding page to display the source-guard binding table for a selected interface.

Parameters

These parameters are displayed:

Query by

- ◆ **Port** – A port on this switch.
- ◆ **VLAN** – ID of a configured VLAN (Range: 1-4094)
- ◆ **MAC Address** – A valid unicast MAC address.
- ◆ **IPv6 Address** – A valid global unicast IPv6 address.

Dynamic Binding List

- ◆ **VLAN** – VLAN to which this entry is bound.
- ◆ **MAC Address** – Physical address associated with the entry.
- ◆ **Interface** – Port to which this entry is bound.
- ◆ **IPv6 Address** – IPv6 address corresponding to the client.
- ◆ **Type** – Shows the entry type:
 - **DHCP** – Dynamic DHCPv6 binding, stateful address.
 - **ND** – Dynamic Neighbor Discovery binding, stateless address.

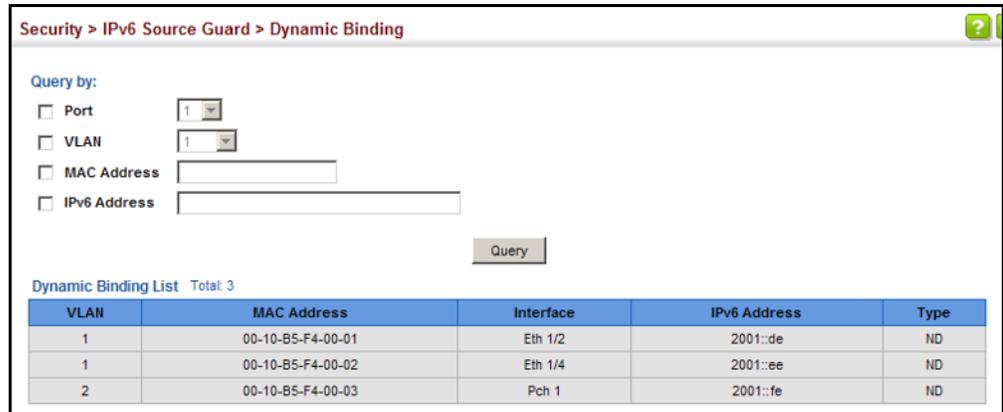
Web Interface

To display the binding table for IPv6 Source Guard:

1. Click Security, IPv6 Source Guard, Dynamic Binding.
2. Mark the search criteria, and enter the required values.

3. Click Query

Figure 187: Showing the IPv6 Source Guard Binding Table



DHCP Snooping

The addresses assigned to DHCP clients on insecure ports can be carefully controlled using the dynamic bindings registered with DHCP Snooping (or using the static bindings configured with IP Source Guard). DHCP snooping allows a switch to protect a network from rogue DHCP servers or other devices which send port-related information to a DHCP server. This information can be useful in tracking an IP address back to a physical port.

Command Usage

DHCP Snooping Process

- ◆ Network traffic may be disrupted when malicious DHCP messages are received from an outside source. DHCP snooping is used to filter DHCP messages received on a non-secure interface from outside the network or fire wall. When DHCP snooping is enabled globally and enabled on a VLAN interface, DHCP messages received on an untrusted interface from a device not listed in the DHCP snooping table will be dropped.
- ◆ Table entries are only learned for trusted interfaces. An entry is added or removed dynamically to the DHCP snooping table when a client receives or releases an IP address from a DHCP server. Each entry includes a MAC address, IP address, lease time, VLAN identifier, and port identifier.
- ◆ The rate limit for the number of DHCP messages that can be processed by the switch is 100 packets per second. Any DHCP packets in excess of this limit are dropped.
- ◆ When DHCP snooping is enabled, DHCP messages entering an untrusted interface are filtered based upon dynamic entries learned via DHCP snooping.

- ◆ Filtering rules are implemented as follows:
 - If the global DHCP snooping is disabled, all DHCP packets are forwarded.
 - If DHCP snooping is enabled globally, and also enabled on the VLAN where the DHCP packet is received, all DHCP packets are forwarded for a *trusted* port. If the received packet is a DHCP ACK message, a dynamic DHCP snooping entry is also added to the binding table.
 - If DHCP snooping is enabled globally, and also enabled on the VLAN where the DHCP packet is received, but the port is *not trusted*, it is processed as follows:
 - If the DHCP packet is a reply packet from a DHCP server (including OFFER, ACK or NAK messages), the packet is dropped.
 - If the DHCP packet is from a client, such as a DECLINE or RELEASE message, the switch forwards the packet only if the corresponding entry is found in the binding table.
 - If the DHCP packet is from a client, such as a DISCOVER, REQUEST, INFORM, DECLINE or RELEASE message, the packet is forwarded if MAC address verification is disabled. However, if MAC address verification is enabled, then the packet will only be forwarded if the client's hardware address stored in the DHCP packet is the same as the source MAC address in the Ethernet header.
 - If the DHCP packet is not a recognizable type, it is dropped.
 - If a DHCP packet from a client passes the filtering criteria above, it will only be forwarded to trusted ports in the same VLAN.
 - If a DHCP packet is from server is received on a trusted port, it will be forwarded to both trusted and untrusted ports in the same VLAN.
 - If the DHCP snooping is globally disabled, all dynamic bindings are removed from the binding table.
 - *Additional considerations when the switch itself is a DHCP client* – The port(s) through which the switch submits a client request to the DHCP server must be configured as trusted. Note that the switch will not add a dynamic entry for itself to the binding table when it receives an ACK message from a DHCP server. Also, when the switch sends out DHCP client packets for itself, no filtering takes place. However, when the switch receives any messages from a DHCP server, any packets received from untrusted ports are dropped.

DHCP Snooping Option 82

- ◆ DHCP provides a relay mechanism for sending information about its DHCP clients or the relay agent itself to the DHCP server. Also known as DHCP Option 82, it allows compatible DHCP servers to use the information when assigning IP addresses, or to set other services or policies for clients. It is also an effective tool in preventing malicious network attacks from attached clients on DHCP services, such as IP Spoofing, Client Identifier Spoofing, MAC Address Spoofing, and Address Exhaustion.
- ◆ DHCP Snooping must be enabled for Option 82 information to be inserted into request packets.
- ◆ When the DHCP Snooping Information Option 82 is enabled, the requesting client (or an intermediate relay agent that has used the information fields to describe itself) can be identified in the DHCP request packets forwarded by the switch and in reply packets sent back from the DHCP server. This information may specify the MAC address or IP address of the requesting device (that is, the switch in this context).

By default, the switch also fills in the Option 82 circuit-id field with information indicating the local interface over which the switch received the DHCP client request, including the port and VLAN ID. This allows DHCP client-server exchange messages to be forwarded between the server and client without having to flood them to the entire VLAN.

- ◆ If DHCP Snooping Information Option 82 is enabled on the switch, information may be inserted into a DHCP request packet received over any VLAN (depending on DHCP snooping filtering rules). The information inserted into the relayed packets includes the circuit-id and remote-id, as well as the gateway Internet address.
- ◆ When the switch receives DHCP packets from clients that already include DHCP Option 82 information, the switch can be configured to set the action policy for these packets. The switch can either drop the DHCP packets, keep the existing information, or replace it with the switch's relay information.

DHCP Snooping Global Configuration Use the IP Service > DHCP > Snooping (Configure Global) page to enable DHCP Snooping globally on the switch, or to configure MAC Address Verification.

Parameters

These parameters are displayed:

General

- ◆ **DHCP Snooping Status** – Enables DHCP snooping globally. (Default: Disabled)
- ◆ **DHCP Snooping MAC-Address Verification** – Enables or disables MAC address verification. If the source MAC address in the Ethernet header of the

packet is not same as the client's hardware address in the DHCP packet, the packet is dropped. (Default: Enabled)

- ◆ **DHCP Snooping Rate Limit** – Sets the maximum number of DHCP packets that can be trapped by the switch for DHCP snooping. (Range: 1-2048 packets/second)

Information

- ◆ **DHCP Snooping Information Option Status** – Enables or disables DHCP Option 82 information relay. (Default: Disabled)
- ◆ **DHCP Snooping Information Option Sub-option Format** – Enables or disables use of sub-type and sub-length fields in circuit-ID (CID) and remote-ID (RID) in Option 82 information. (Default: Disabled)
- ◆ **DHCP Snooping Information Option Remote ID** – Specifies the MAC address, IP address, or arbitrary identifier of the requesting device (i.e., the switch in this context).
 - **MAC Address** – Inserts a MAC address in the remote ID sub-option for the DHCP snooping agent (i.e., the MAC address of the switch's CPU). This attribute can be encoded in Hexadecimal or ASCII.
 - **IP Address** – Inserts an IP address in the remote ID sub-option for the DHCP snooping agent (i.e., the IP address of the management interface). This attribute can be encoded in Hexadecimal or ASCII.
 - *string* - An arbitrary string inserted into the remote identifier field. (Range: 1-32 characters)
- ◆ **DHCP Snooping Information Option Policy** – Specifies how to handle DHCP client request packets which already contain Option 82 information.
 - **Drop** – Drops the client's request packet instead of relaying it.
 - **Keep** – Retains the Option 82 information in the client request, and forwards the packets to trusted ports.
 - **Replace** – Replaces the Option 82 information circuit-id and remote-id fields in the client's request with information about the relay agent itself, inserts the relay agent's address (when DHCP snooping is enabled), and forwards the packets to trusted ports. (This is the default policy.)

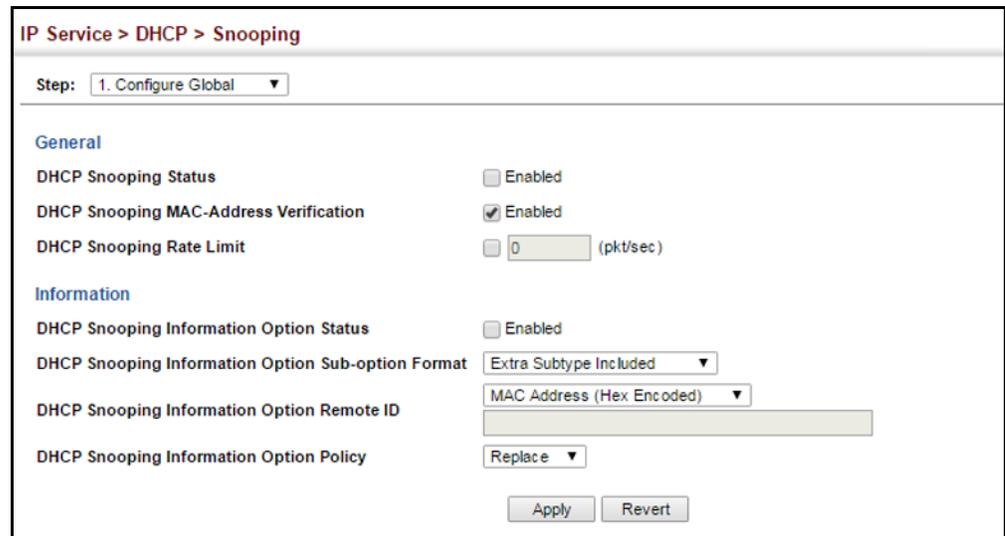
Web Interface

To configure global settings for DHCP Snooping:

1. Click IP Service, DHCP, Snooping.
2. Select Configure Global from the Step list.
3. Select the required options for the general DHCP snooping process and for the DHCP Option 82 information.

4. Click Apply

Figure 188: Configuring Global Settings for DHCP Snooping



DHCP Snooping VLAN Configuration Use the IP Service > DHCP > Snooping (Configure VLAN) page to enable or disable DHCP snooping on specific VLANs.

Command Usage

- ◆ When DHCP snooping is enabled globally on the switch, and enabled on the specified VLAN, DHCP packet filtering will be performed on any untrusted ports within the VLAN.
- ◆ When the DHCP snooping is globally disabled, DHCP snooping can still be configured for specific VLANs, but the changes will not take effect until DHCP snooping is globally re-enabled.
- ◆ When DHCP snooping is globally enabled, and DHCP snooping is then disabled on a VLAN, all dynamic bindings learned for this VLAN are removed from the binding table.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of a configured VLAN. (Range: 1-4094)
- ◆ **DHCP Snooping Status** – Enables or disables DHCP snooping for the selected VLAN. When DHCP snooping is enabled globally on the switch, and enabled on the specified VLAN, DHCP packet filtering will be performed on any untrusted ports within the VLAN. (Default: Disabled)

Web Interface

To configure global settings for DHCP Snooping:

1. Click IP Service, DHCP, Snooping.
2. Select Configure VLAN from the Step list.
3. Enable DHCP Snooping on any existing VLAN.
4. Click Apply

Figure 189: Configuring DHCP Snooping on a VLAN

The screenshot shows a web interface for configuring DHCP Snooping. The breadcrumb path is 'IP Service > DHCP > Snooping'. Below the breadcrumb, there is a 'Step:' dropdown menu currently showing '2. Configure VLAN'. Underneath, there is a 'VLAN' dropdown menu showing '1'. Below that, there is a 'DHCP Snooping Status' section with a checked checkbox and the text 'Enabled'. At the bottom right of the form, there are two buttons: 'Apply' and 'Revert'.

Configuring Interfaces for DHCP Snooping

Use the IP Service > DHCP > Snooping (Configure Interface) page to configure switch interfaces as trusted or untrusted.

Command Usage

- ◆ A trusted interface is an interface that is configured to receive only messages from within the network. An untrusted interface is an interface that is configured to receive messages from outside the network or fire wall.
- ◆ When DHCP snooping is enabled both globally and on a VLAN, DHCP packet filtering will be performed on any untrusted interface within the VLAN.
- ◆ When an untrusted interface is changed to a trusted interface, all the dynamic DHCP snooping bindings associated with this interface are removed.
- ◆ Set all interfaces connected to DHCP servers within the local network or fire wall to trusted state. Set all other interfaces outside the local network or fire wall to untrusted state.

Parameters

These parameters are displayed:

- ◆ **Interface**
 - Port identifier. (Range: 1-52)
 - Trunk identifier. (Range: 1-26)
- ◆ **Trust Status** – Enables or disables a port as trusted. (Default: Disabled)

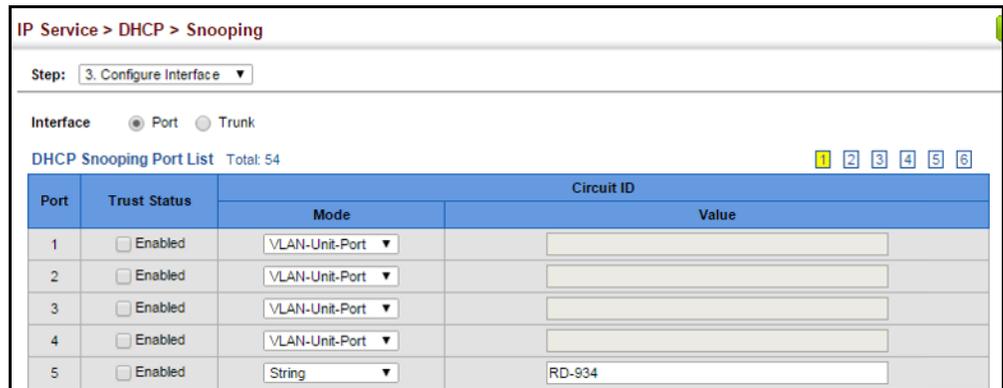
- ◆ **Circuit ID** – Specifies DHCP Option 82 circuit ID suboption information.
 - **Mode** – Specifies the default string “VLAN-Unit-Port” or an arbitrary string. (Default: VLAN-Unit-Port)
 - **Value** – An arbitrary string inserted into the circuit identifier field. (Range: 1-32 characters)

Web Interface

To configure global settings for DHCP Snooping:

1. Click IP Service, DHCP, Snooping.
2. Select Configure Interface from the Step list.
3. Set any ports within the local network or firewall to trusted.
4. Specify the mode used for sending circuit ID information, and an arbitrary string if required.
5. Click Apply

Figure 190: Configuring the Port Mode for DHCP Snooping



Displaying DHCP Snooping Binding Information

Use the IP Service > DHCP > Snooping (Show Information) page to display entries in the binding table.

Parameters

These parameters are displayed:

- ◆ **MAC Address** – Physical address associated with the entry.
- ◆ **IP Address** – IP address corresponding to the client.
- ◆ **Lease Time** – The time for which this IP address is leased to the client.

- ◆ **Type** – Entry types include:
 - **DHCP-Snooping** – Dynamically snooped.
 - **Static-DHCP-SNP** – Statically configured.
- ◆ **VLAN** – VLAN to which this entry is bound.
- ◆ **Interface** – Port or trunk to which this entry is bound.
- ◆ **Store** – Writes all dynamically learned snooping entries to flash memory. This function can be used to store the currently learned dynamic DHCP snooping entries to flash memory. These entries will be restored to the snooping table when the switch is reset. However, note that the lease time shown for a dynamic entry that has been restored from flash memory will no longer be valid.
- ◆ **Clear** – Removes all dynamically learned snooping entries from flash memory.

Web Interface

To display the binding table for DHCP Snooping:

1. Click IP Service, DHCP, Snooping.
2. Select Show Information from the Step list.
3. Use the Store or Clear function if required.

Figure 191: Displaying the Binding Table for DHCP Snooping

IP Service > DHCP > Snooping

Step: 4. Show Information

DHCP Snooping Binding List Total: 6

MAC Address	IP Address	Lease Time (seconds)	Type	VLAN	Interface
00-10-B5-F4-00-01	10.2.44.96	5	DHCP-Snooping	1	Trunk 1
00-10-B5-F4-00-02	10.3.44.96	15	Static-DHCP-SNP	1	Unit 1 / Port 2
00-10-B5-F4-00-03	10.4.44.96	25	DHCP-Snooping	1	Unit 1 / Port 3
00-10-B5-F4-00-04	10.5.44.96	10	Static-DHCP-SNP	1	Trunk 4
00-10-B5-F4-00-05	10.6.44.96	10	DHCP-Snooping	1	Unit 1 / Port 5
00-10-B5-F4-00-06	10.7.44.96	5	Static-DHCP-SNP	1	Unit 1 / Port 6

Store Click the button to Store DHCP Snooping binding entries to flash.

Clear Click the button to Clear DHCP Snooping binding entries from flash.

Basic Administration Protocols

This chapter describes basic administration tasks including:

- ◆ **Event Logging** – Sets conditions for logging event messages to system memory or flash memory, configures conditions for sending trap messages to remote log servers, and configures trap reporting to remote hosts using Simple Mail Transfer Protocol (SMTP).
- ◆ **Link Layer Discovery Protocol (LLDP)** – Configures advertisement of basic information about the local switch, or discovery of information about neighboring devices on the local broadcast domain.
- ◆ **Simple Network Management Protocol (SNMP)** – Configures switch management through SNMPv1, SNMPv2c or SNMPv3.
- ◆ **Remote Monitoring (RMON)** – Configures local collection of detailed statistics or events which can be subsequently retrieved through SNMP.
- ◆ **Connectivity Fault Management (CFM)** – This protocol provides proactive connectivity monitoring using continuity check messages, fault verification through loop back messages, and fault isolation by examining end-to-end connections between provider edge devices or between customer edge devices.
- ◆ **UniDirectional Link Detection (UDLD)** – Detects and disables unidirectional Ethernet fiber or copper links.

Configuring Event Logging

The switch allows you to control the logging of error messages, including the type of events that are recorded in switch memory, logging to a remote System Log (syslog) server, and displays a list of recent event messages.

System Log Configuration Use the Administration > Log > System (Configure Global) page to enable or disable event logging, and specify which levels are logged to RAM or flash memory.

Severe error messages that are logged to flash memory are permanently stored in the switch to assist in troubleshooting network problems. Up to 4096 log entries can be stored in the flash memory, with the oldest entries being overwritten first when the available log memory (256 kilobytes) has been exceeded.

The System Logs page allows you to configure and limit system messages that are logged to flash or RAM memory. The default is for event levels 0 to 3 to be logged to flash and levels 0 to 7 to be logged to RAM.

Parameters

These parameters are displayed:

- ◆ **System Log Status** – Enables/disables the logging of debug or error messages to the logging process. (Default: Enabled)
- ◆ **Flash Level** – Limits log messages saved to the switch’s permanent flash memory for all levels up to the specified level. For example, if level 3 is specified, all messages from level 0 to level 3 will be logged to flash. (Range: 0-7, Default: 3)

Table 26: Logging Levels

Level	Severity Name	Description
7	Debug	Debugging messages
6	Informational	Informational messages only
5	Notice	Normal but significant condition, such as cold start
4	Warning	Warning conditions (e.g., return false, unexpected return)
3	Error	Error conditions (e.g., invalid input, default used)
2	Critical	Critical conditions (e.g., memory allocation, or free memory error - resource exhausted)
1	Alert	Immediate action needed
0	Emergency	System unusable

* There are only Level 2, 5 and 6 error messages for the current firmware release.

- ◆ **RAM Level** – Limits log messages saved to the switch’s temporary RAM memory for all levels up to the specified level. For example, if level 7 is specified, all messages from level 0 to level 7 will be logged to RAM. (Range: 0-7, Default: 7)



Note: The Flash Level must be equal to or less than the RAM Level.

Note: All log messages are retained in RAM and Flash after a warm restart (i.e., power is reset through the command interface).

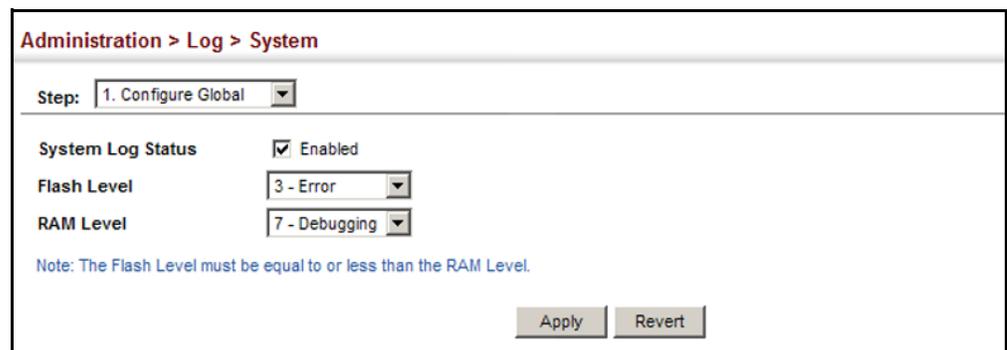
Note: All log messages are retained in Flash and purged from RAM after a cold restart (i.e., power is turned off and then on through the power source).

Web Interface

To configure the logging of error messages to system memory:

1. Click Administration, Log, System.
2. Select Configure Global from the Step list.
3. Enable or disable system logging, set the level of event messages to be logged to flash memory and RAM.
4. Click Apply.

Figure 193: Configuring Settings for System Memory Logs



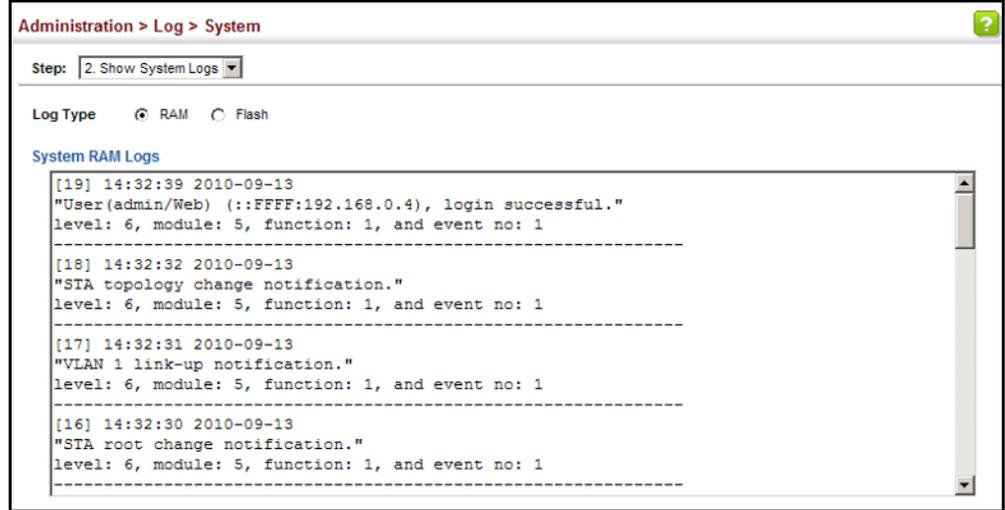
The screenshot shows a web interface for configuring system memory logs. The breadcrumb navigation is "Administration > Log > System". The "Step:" dropdown menu is set to "1. Configure Global". The "System Log Status" is checked and set to "Enabled". The "Flash Level" dropdown menu is set to "3 - Error". The "RAM Level" dropdown menu is set to "7 - Debugging". A note below the settings states: "Note: The Flash Level must be equal to or less than the RAM Level." At the bottom right, there are two buttons: "Apply" and "Revert".

To show the error messages logged to system or flash memory:

1. Click Administration, Log, System.
2. Select Show System Logs from the Step list.
3. Click RAM to display log messages stored in system memory, or Flash to display messages stored in flash memory.

This page allows you to scroll through the logged system and event messages. The switch can store up to 2048 log entries in temporary random access memory (RAM; i.e., memory flushed on power reset) and up to 4096 entries in permanent flash memory.

Figure 194: Showing Error Messages Logged to System Memory



Remote Log Configuration Use the Administration > Log > Remote page to send log messages to syslog servers or other management stations. You can also limit the event messages sent to only those messages below a specified level.

Parameters

These parameters are displayed:

- ◆ **Remote Log Status** – Enables/disables the logging of debug or error messages to the remote logging process. (Default: Disabled)
- ◆ **Logging Facility** – Sets the facility type for remote logging of syslog messages. There are eight facility types specified by values of 16 to 23. The facility type is used by the syslog server to dispatch log messages to an appropriate service.

The attribute specifies the facility type tag sent in syslog messages (see RFC 3164). This type has no effect on the kind of messages reported by the switch. However, it may be used by the syslog server to process messages, such as sorting or storing messages in the corresponding database. (Range: 16-23, Default: 23)
- ◆ **Logging Trap Level** – Limits log messages that are sent to the remote syslog server for all levels up to the specified level. For example, if level 3 is specified, all messages from level 0 to level 3 will be sent to the remote server. (Range: 0-7, Default: 7)
- ◆ **Server IP Address** – Specifies the IPv4 or IPv6 address of a remote server which will be sent syslog messages.
- ◆ **Port** - Specifies the UDP port number used by the remote server. (Range: 1-65535; Default: 514)

Web Interface

To configure the logging of error messages to remote servers:

1. Click Administration, Log, Remote.
2. Enable remote logging, specify the facility type to use for the syslog messages, and enter the IP address of the remote servers.
3. Click Apply.

Figure 195: Configuring Settings for Remote Logging of Error Messages

The screenshot shows a web interface titled "Administration > Log > Remote". It contains the following configuration options:

- Remote Log Status:** A checkbox labeled "Enabled" which is currently unchecked.
- Logging Facility:** A dropdown menu with "23 - Local use 7" selected.
- Logging Trap Level:** A dropdown menu with "0 - System unusable" selected.
- Server IP Address 1:** A text input field containing "192.168.0.4".
- Server IP Address 2:** An empty text input field.
- Server IP Address 3:** An empty text input field.
- Server IP Address 4:** An empty text input field.
- Server IP Address 5:** An empty text input field.

Each IP address field is paired with a "Port" field. The port for the first server is "514", while the others are empty. At the bottom of the form are "Apply" and "Revert" buttons.

Link Layer Discovery Protocol

Link Layer Discovery Protocol (LLDP) is used to discover basic information about neighboring devices on the local broadcast domain. LLDP is a Layer 2 protocol that uses periodic broadcasts to advertise information about the sending device. Advertised information is represented in Type Length Value (TLV) format according to the IEEE 802.1AB standard, and can include details such as device identification, capabilities and configuration settings. LLDP also defines how to store and maintain information gathered about the neighboring network nodes it discovers.

Link Layer Discovery Protocol - Media Endpoint Discovery (LLDP-MED) is an extension of LLDP intended for managing endpoint devices such as Voice over IP phones and network switches. The LLDP-MED TLVs advertise information such as network policy, power, inventory, and device location details. LLDP and LLDP-MED information can be used by SNMP applications to simplify troubleshooting, enhance network management, and maintain an accurate network topology.

Setting LLDP Timing Attributes Use the Administration > LLDP (Configure Global) page to set attributes for general functions such as globally enabling LLDP on the switch, setting the message ageout time, and setting the frequency for broadcasting general advertisements or reports about changes in the LLDP MIB.

Parameters

These parameters are displayed:

- ◆ **LLDP** – Enables LLDP globally on the switch. (Default: Enabled)
- ◆ **Transmission Interval** – Configures the periodic transmit interval for LLDP advertisements. (Range: 5-32768 seconds; Default: 30 seconds)
- ◆ **Hold Time Multiplier** – Configures the time-to-live (TTL) value sent in LLDP advertisements as shown in the formula below. (Range: 2-10; Default: 4)

The time-to-live tells the receiving LLDP agent how long to retain all information pertaining to the sending LLDP agent if it does not transmit updates in a timely manner. Note that if the local interface attached to a remote device is shut down or otherwise disabled, information about the remote device is purged immediately.

TTL in seconds is based on the following rule:
minimum value ((Transmission Interval * Holdtime Multiplier), or 65535)

Therefore, the default TTL is $4 * 30 = 120$ seconds.

- ◆ **Delay Interval** – Configures a delay between the successive transmission of advertisements initiated by a change in local LLDP MIB variables. (Range: 1-8192 seconds; Default: 2 seconds)

The transmit delay is used to prevent a series of successive LLDP transmissions during a short period of rapid changes in local LLDP MIB objects, and to increase the probability that multiple, rather than single changes, are reported in each transmission.

This attribute must comply with the rule:
 $(4 * \text{Delay Interval}) \leq \text{Transmission Interval}$
- ◆ **Reinitialization Delay** – Configures the delay before attempting to re-initialize after LLDP ports are disabled or the link goes down. (Range: 1-10 seconds; Default: 2 seconds)

When LLDP is re-initialized on a port, all information in the remote systems LLDP MIB associated with this port is deleted.
- ◆ **Notification Interval** – Configures the allowed interval for sending SNMP notifications about LLDP MIB changes. (Range: 5-3600 seconds; Default: 5 seconds)

This parameter only applies to SNMP applications which use data stored in the LLDP MIB for network monitoring or management.

Information about changes in LLDP neighbors that occur between SNMP notifications is not transmitted. Only state changes that exist at the time of a

notification are included in the transmission. An SNMP agent should therefore periodically check the value of `IldpStatsRemTableLastChangeTime` to detect any `IldpRemTablesChange` notification-events missed due to throttling or transmission loss.

- ◆ **MED Fast Start Count** – Configures the amount of LLDP MED Fast Start LLDPDUs to transmit during the activation process of the LLDP-MED Fast Start mechanism. (Range: 1-10 packets; Default: 4 packets)

The MED Fast Start Count parameter is part of the timer which ensures that the LLDP-MED Fast Start mechanism is active for the port. LLDP-MED Fast Start is critical to the timely startup of LLDP, and therefore integral to the rapid availability of Emergency Call Service.

Web Interface

To configure LLDP timing attributes:

1. Click Administration, LLDP.
2. Select Configure Global from the Step list.
3. Enable LLDP, and modify any of the timing parameters as required.
4. Click Apply.

Figure 196: Configuring LLDP Timing Attributes

The screenshot shows the 'Administration > LLDP' configuration page. At the top, there is a 'Step:' dropdown menu set to '1. Configure Global'. Below this, the 'LLDP' section is checked as 'Enabled'. The following parameters are listed with their respective values in input boxes and units:

Parameter	Value	Unit
Transmission Interval (5-32768)	30	sec
Hold Time Multiplier (2-10)	4	
Delay Interval (1-8192)	2	sec
Reinitialization Delay (1-10)	2	sec
Notification Interval (5-3600)	5	sec
MED Fast Start Count (1-10)	4	

Below the table, a note states: 'Note: The Transmission Interval must be greater than or equal to 4 times the Delay Interval.' At the bottom right, there are two buttons: 'Apply' and 'Revert'.

Configuring LLDP Interface Attributes

Use the Administration > LLDP (Configure Interface – Configure General) page to specify the message attributes for individual interfaces, including whether messages are transmitted, received, or both transmitted and received, whether SNMP notifications are sent, and the type of information advertised.

Parameters

These parameters are displayed:

- ◆ **Admin Status** – Enables LLDP message transmit and receive modes for LLDP Protocol Data Units. (Options: Tx only, Rx only, TxRx, Disabled; Default: TxRx)

- ◆ **SNMP Notification** – Enables the transmission of SNMP trap notifications about LLDP and LLDP-MED changes. (Default: Enabled)

This option sends out SNMP trap notifications to designated target stations at the interval specified by the Notification Interval in the preceding section. Trap notifications include information about state changes in the LLDP MIB (IEEE 802.1AB), the LLDP-MED MIB (ANSI/TIA-1057), or vendor-specific LLDP-EXT-DOT1 and LLDP-EXT-DOT3 MIBs.

For information on defining SNMP trap destinations, see [“Specifying Trap Managers” on page 372](#).

Information about additional changes in LLDP neighbors that occur between SNMP notifications is not transmitted. Only state changes that exist at the time of a trap notification are included in the transmission. An SNMP agent should therefore periodically check the value of `IldpStatsRemTableLastChangeTime` to detect any `IldpRemTablesChange` notification-events missed due to throttling or transmission loss.

- ◆ **MED Notification** – Enables the transmission of SNMP trap notifications about LLDP-MED changes in remote neighbors. (Default: Disabled)
- ◆ **Basic Optional TLVs** – Configures basic information included in the TLV field of advertised messages.

- **Management Address** – The management address protocol packet includes the IPv4 address of the switch. If no management address is available, the address should be the MAC address for the CPU or for the port sending this advertisement. (Default: Enabled)

The management address TLV may also include information about the specific interface associated with this address, and an object identifier indicating the type of hardware component or protocol entity associated with this address. The interface number and OID are included to assist SNMP applications in the performance of network discovery by indicating enterprise specific or other starting points for the search, such as the Interface or Entity MIB.

Since there are typically a number of different addresses associated with a Layer 3 device, an individual LLDP PDU may contain more than one management address TLV.

Every management address TLV that reports an address that is accessible on a port and protocol VLAN through the particular port should be accompanied by a port and protocol VLAN TLV that indicates the VLAN identifier (VID) associated with the management address reported by this TLV.

- **Port Description** – The port description is taken from the ifDescr object in RFC 2863, which includes information about the manufacturer, the product name, and the version of the interface hardware/software. (Default: Enabled)
 - **System Capabilities** – The system capabilities identifies the primary function(s) of the system and whether or not these primary functions are enabled. The information advertised by this TLV is described in IEEE 802.1AB. (Default: Enabled)
 - **System Description** – The system description is taken from the sysDescr object in RFC 3418, which includes the full name and version identification of the system's hardware type, software operating system, and networking software. (Default: Enabled)
 - **System Name** – The system name is taken from the sysName object in RFC 3418, which contains the system's administratively assigned name. To configure the system name, see [“Displaying System Information” on page 72](#). (Default: Enabled)
- ◆ **802.1 Organizationally Specific TLVs** – Configures IEEE 802.1 information included in the TLV field of advertised messages.
 - **Protocol Identity** – The protocols that are accessible through this interface. (Default: Enabled)
 - **VLAN ID** – The port's default VLAN identifier (PVID) indicates the VLAN with which untagged or priority-tagged frames are associated (see [“IEEE 802.1Q VLANs” on page 145](#)). (Default: Enabled)
 - **VLAN Name** – The name of all VLANs to which this interface has been assigned (see [“IEEE 802.1Q VLANs” on page 145](#)). (Default: Enabled)
 - **Port and Protocol VLAN ID** – The port-based protocol VLANs configured on this interface. (Default: Enabled)
- ◆ **802.3 Organizationally Specific TLVs** – Configures IEEE 802.3 information included in the TLV field of advertised messages.
 - **Link Aggregation** – The link aggregation capabilities, aggregation status of the link, and the IEEE 802.3 aggregated port identifier if this interface is currently a link aggregation member. (Default: Enabled)

- **Max Frame Size** – The maximum frame size. (See “Configuring Support for Jumbo Frames” on page 74 for information on configuring the maximum frame size for this switch (Default: Enabled)
- **MAC/PHY Configuration/Status** – The MAC/PHY configuration and status which includes information about auto-negotiation support/capabilities, and operational Multistation Access Unit (MAU) type. (Default: Enabled)
- ◆ **MED TLVs** – Configures general information included in the MED TLV field of advertised messages.
 - **Capabilities** – This option advertises LLDP-MED TLV capabilities, allowing Media Endpoint and Connectivity Devices to efficiently discover which LLDP-MED related TLVs are supported on the switch. (Default: Enabled)
 - **Inventory** – This option advertises device details useful for inventory management, such as manufacturer, model, software version and other pertinent information. (Default: Enabled)
 - **Location** – This option advertises location identification details. (Default: Enabled)
 - **Network Policy** – This option advertises network policy configuration information, aiding in the discovery and diagnosis of VLAN configuration mismatches on a port. Improper network policy configurations frequently result in voice quality degradation or complete service disruption. (Default: Enabled)
- ◆ **MED-Location Civic Address** – Configures information for the location of the attached device included in the MED TLV field of advertised messages, including the country and the device type.
 - **Country** – The two-letter ISO 3166 country code in capital ASCII letters. (Example: DK, DE or US)
 - **Device entry refers to** – The type of device to which the location applies:
 - Location of DHCP server.
 - Location of network element closest to client.
 - Location of client. (This is the default.)

Web Interface

To configure LLDP interface attributes:

1. Click Administration, LLDP.
2. Select Configure Interface from the Step list.
3. Select Configure General from the Action list.

4. Select an interface from the Port or Trunk list.
5. Set the LLDP transmit/receive mode, specify whether or not to send SNMP trap messages, and select the information to advertise in LLDP messages.
6. Click Apply.

Figure 197: Configuring LLDP Interface Attributes

The screenshot shows the 'Administration > LLDP' configuration page. At the top, there are dropdowns for 'Step: 2. Configure Interface' and 'Action: Configure General'. Below this, the 'Interface' section has radio buttons for 'Port' (selected) and 'Trunk', with a dropdown for '24'. The 'Admin Status' is set to 'Tx Rx'. 'SNMP Notification' and 'MED Notification' are both checked and set to 'Enabled'. Under 'Basic Optional TLVs', several checkboxes are checked: Management Address, Port Description, System Capabilities, System Description, and System Name. Under '802.1 Organizationally Specific TLVs', Protocol Identity, VLAN ID, VLAN Name, and Port and Protocol VLAN ID are checked. Under '802.3 Organizationally Specific TLVs', Link Aggregation, Max Frame Size, and MAC/PHY Configuration/Status are checked. Under 'MED TLVs', Capabilities, Inventory, Location, and Network Policy are checked. The 'MED-Location Civic Address' section has a 'Country' dropdown set to 'US' and a 'Device entry refers to' dropdown set to 'Location of the client'. A note at the bottom states: 'Note: The country string shall be a two-letter ISO 3166 country code, e.g. US'. At the bottom right, there are 'Apply' and 'Revert' buttons.

Configuring LLDP Interface Civic-Address

Use the Administration > LLDP (Configure Interface – Add CA-Type) page to specify the physical location of the device attached to an interface.

Command Usage

- ◆ Use the Civic Address type (CA-Type) to advertise the physical location of the device attached to an interface, including items such as the city, street number, building and room information. The address location is specified as a type and value pair, with the civic address type defined in RFC 4776. The following table describes some of the CA type numbers and provides examples.

Table 27: LLDP MED Location CA Types

CA Type	Description	CA Value Example
1	National subdivisions (state, canton, province)	California
2	County, parish	Orange
3	City, township	Irvine
4	City division, borough, city district	West Irvine

Table 27: LLDP MED Location CA Types (Continued)

CA Type	Description	CA Value Example
5	Neighborhood, block	Riverside
6	Group of streets below the neighborhood level	Exchange
18	Street suffix or type	Avenue
19	House number	320
20	House number suffix	A
21	Landmark or vanity address	Tech Center
26	Unit (apartment, suite)	Apt 519
27	Floor	5
28	Room	509B

- ◆ Any number of CA type and value pairs can be specified for the civic address location, as long as the total does not exceed 250 characters.

Parameters

These parameters are displayed:

- ◆ **CA-Type** – Descriptor of the data civic address value. (Range: 0-255)
- ◆ **CA-Value** – Description of a location. (Range: 1-32 characters)

Web Interface

To specify the physical location of the attached device:

1. Click Administration, LLDP.
2. Select Configure Interface from the Step list.
3. Select Add CA-Type from the Action list.
4. Select an interface from the Port or Trunk list.
5. Specify a CA-Type and CA-Value pair.
6. Click Apply.

Figure 198: Configuring the Civic Address for an LLDP Interface

To show the physical location of the attached device:

1. Click Administration, LLDP.
2. Select Configure Interface from the Step list.
3. Select Show CA-Type from the Action list.
4. Select an interface from the Port or Trunk list.

Figure 199: Showing the Civic Address for an LLDP Interface

CA-Type	CA-Value
1	California

Displaying LLDP Local Device Information

Use the Administration > LLDP (Show Local Device Information) page to display information about the switch, such as its MAC address, chassis ID, management IP address, and port information.

Parameters

These parameters are displayed:

General Settings

- ◆ **Chassis Type** – Identifies the chassis containing the IEEE 802 LAN entity associated with the transmitting LLDP agent. There are several ways in which a chassis may be identified and a chassis ID subtype is used to indicate the type of component being referenced by the chassis ID field.

Table 28: Chassis ID Subtype

ID Basis	Reference
Chassis component	EntPhysicalAlias when entPhysClass has a value of 'chassis(3)' (IETF RFC 2737)
Interface alias	IfAlias (IETF RFC 2863)
Port component	EntPhysicalAlias when entPhysicalClass has a value 'port(10)' or 'backplane(4)' (IETF RFC 2737)
MAC address	MAC address (IEEE Std 802-2001)
Network address	networkAddress
Interface name	ifName (IETF RFC 2863)
Locally assigned	locally assigned

- ◆ **Chassis ID** – An octet string indicating the specific identifier for the particular chassis in this system.
- ◆ **System Name** – A string that indicates the system’s administratively assigned name (see [“Displaying System Information” on page 72](#)).
- ◆ **System Description** – A textual description of the network entity. This field is also displayed by the **show system** command.
- ◆ **System Capabilities Supported** – The capabilities that define the primary function(s) of the system.

Table 29: System Capabilities

ID Basis	Reference
Other	—
Repeater	IETF RFC 2108
Bridge	IETF RFC 2674
WLAN Access Point	IEEE 802.11 MIB
Router	IETF RFC 1812
Telephone	IETF RFC 2011
DOCSIS cable device	IETF RFC 2669 and IETF RFC 2670
End Station Only	IETF RFC 2011

- ◆ **System Capabilities Enabled** – The primary function(s) of the system which are currently enabled. Refer to the preceding table.
- ◆ **Management Address** – The management address associated with the local system. If no management address is available, the address should be the MAC address for the CPU or for the port sending this advertisement.

Interface Settings

The attributes listed below apply to both port and trunk interface types. When a trunk is listed, the descriptions apply to the first port of the trunk.

- ◆ **Port/Trunk Description** – A string that indicates the port or trunk description. If RFC 2863 is implemented, the ifDescr object should be used for this field.
- ◆ **Port/Trunk ID** – A string that contains the specific identifier for the port or trunk from which this LLDPDU was transmitted.

Interface Details

The attributes listed below apply to both port and trunk interface types. When a trunk is listed, the descriptions apply to the first port of the trunk.

- ◆ **Local Port/Trunk** – Local interface on this switch.
- ◆ **Port/Trunk ID Type** – There are several ways in which a port may be identified. A port ID subtype is used to indicate how the port is being referenced in the Port ID TLV.

Table 30: Port ID Subtype

ID Basis	Reference
Interface alias	IfAlias (IETF RFC 2863)
Chassis component	EntPhysicalAlias when entPhysClass has a value of 'chassis(3)' (IETF RFC 2737)
Port component	EntPhysicalAlias when entPhysicalClass has a value 'port(10)' or 'backplane(4)' (IETF RFC 2737)
MAC address	MAC address (IEEE Std 802-2001)
Network address	networkAddress
Interface name	ifName (IETF RFC 2863)
Agent circuit ID	agent circuit ID (IETF RFC 3046)
Locally assigned	locally assigned

- ◆ **Port/Trunk ID** – A string that contains the specific identifier for the local interface based on interface subtype used by this switch.
- ◆ **Port/Trunk Description** – A string that indicates the port or trunk description. If RFC 2863 is implemented, the ifDescr object should be used for this field.
- ◆ **MED Capability** – The supported set of capabilities that define the primary function(s) of the interface:
 - LLDP-MED Capabilities
 - Network Policy
 - Location Identification

- Extended Power via MDI – PSE
- Extended Power via MDI – PD
- Inventory

Web Interface

To display LLDP information for the local device:

1. Click Administration, LLDP.
2. Select Show Local Device Information from the Step list.
3. Select General, Port, Port Details, Trunk, or Trunk Details.

Figure 200: Displaying Local Device Information for LLDP (General)

The screenshot shows the 'Administration > LLDP' page. The 'Step' dropdown is set to '3. Show Local Device Information'. Below this, there are radio buttons for 'General' (selected), 'Port', 'Port Details', 'Trunk', and 'Trunk Details'. The main content area is titled 'LLDP Local Device Information' and contains the following details:

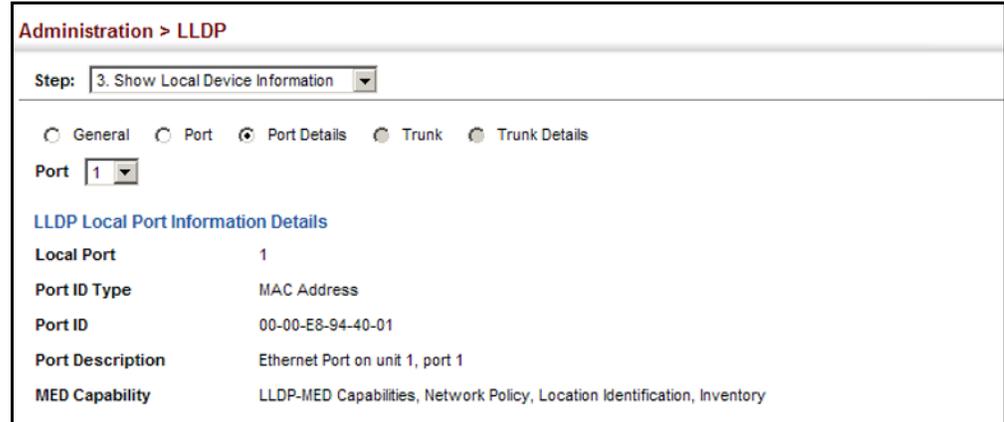
Chassis Type	MAC Address
Chassis ID	00-E0-0C-02-00-FD
System Name	
System Description	ECS5610-52S
System Capabilities Supported	Bridge, Router
System Capabilities Enabled	Bridge, Router
Management Address	192.168.0.2 (IPv4)

Figure 201: Displaying Local Device Information for LLDP (Port)

The screenshot shows the 'Administration > LLDP' page with the 'Port' radio button selected. The 'Step' dropdown is '3. Show Local Device Information'. Below the radio buttons, there is a table titled 'LLDP Local Device Port List' with a total of 54 entries. The table has three columns: 'Port', 'Port Description', and 'Port ID'. The first five rows are visible:

Port	Port Description	Port ID
1	Ethernet Port on unit 1, port 1	70-72-CF-EA-1B-72
2	Ethernet Port on unit 1, port 2	70-72-CF-EA-1B-73
3	Ethernet Port on unit 1, port 3	70-72-CF-EA-1B-74
4	Ethernet Port on unit 1, port 4	70-72-CF-EA-1B-75
5	Ethernet Port on unit 1, port 5	70-72-CF-EA-1B-76

Figure 202: Displaying Local Device Information for LLDP (Port Details)



Displaying LLDP Remote Device Information

Use the Administration > LLDP (Show Remote Device Information) page to display information about devices connected directly to the switch's ports which are advertising information through LLDP, or to display detailed information about an LLDP-enabled device connected to a specific port on the local switch.

Parameters

These parameters are displayed:

Port

- ◆ **Local Port** – The local port to which a remote LLDP-capable device is attached.
- ◆ **Chassis ID** – An octet string indicating the specific identifier for the particular chassis in this system.
- ◆ **Port ID** – A string that contains the specific identifier for the port from which this LLDPDU was transmitted.
- ◆ **System Name** – A string that indicates the system's administratively assigned name.

Port Details

- ◆ **Port** – Port identifier on local switch.
- ◆ **Remote Index** – Index of remote device attached to this port.
- ◆ **Local Port** – The local port to which a remote LLDP-capable device is attached.
- ◆ **Chassis Type** – Identifies the chassis containing the IEEE 802 LAN entity associated with the transmitting LLDP agent. There are several ways in which a chassis may be identified and a chassis ID subtype is used to indicate the type of component being referenced by the chassis ID field. (See [Table 28, "Chassis ID Subtype,"](#) on page 340.)

- ◆ **Chassis ID** – An octet string indicating the specific identifier for the particular chassis in this system.
- ◆ **System Name** – A string that indicates the system’s assigned name.
- ◆ **System Description** – A textual description of the network entity.
- ◆ **Port Type** – Indicates the basis for the identifier that is listed in the Port ID field. See [Table 30, “Port ID Subtype,” on page 341](#).
- ◆ **Port Description** – A string that indicates the port’s description. If RFC 2863 is implemented, the ifDescr object should be used for this field.
- ◆ **Port ID** – A string that contains the specific identifier for the port from which this LLDPDU was transmitted.
- ◆ **System Capabilities Supported** – The capabilities that define the primary function(s) of the system. (See [Table 29, “System Capabilities,” on page 340](#).)
- ◆ **System Capabilities Enabled** – The primary function(s) of the system which are currently enabled. (See [Table 29, “System Capabilities,” on page 340](#).)
- ◆ **Management Address List** – The management addresses for this device. Since there are typically a number of different addresses associated with a Layer 3 device, an individual LLDP PDU may contain more than one management address TLV.

If no management address is available, the address should be the MAC address for the CPU or for the port sending this advertisement.

Port Details – 802.1 Extension Information

- ◆ **Remote Port VID** – The port’s default VLAN identifier (PVID) indicates the VLAN with which untagged or priority-tagged frames are associated.
- ◆ **Remote Port-Protocol VLAN List** – The port-based protocol VLANs configured on this interface, whether the given port (associated with the remote system) supports port-based protocol VLANs, and whether the port-based protocol VLANs are enabled on the given port associated with the remote system.
- ◆ **Remote VLAN Name List** – VLAN names associated with a port.
- ◆ **Remote Protocol Identity List** – Information about particular protocols that are accessible through a port. This object represents an arbitrary local integer value used by this agent to identify a particular protocol identity, and an octet string used to identify the protocols associated with a port of the remote system.

Port Details – 802.3 Extension Port Information

- ◆ **Remote Port Auto-Neg Supported** – Shows whether the given port (associated with remote system) supports auto-negotiation.
- ◆ **Remote Port Auto-Neg Adv-Capability** – The value (bitmap) of the ifMauAutoNegCapAdvertisedBits object (defined in IETF RFC 3636) which is associated with a port on the remote system.

Table 31: Remote Port Auto-Negotiation Advertised Capability

Bit	Capability
0	other or unknown
1	10BASE-T half duplex mode
2	10BASE-T full duplex mode
3	100BASE-T4
4	100BASE-TX half duplex mode
5	100BASE-TX full duplex mode
6	100BASE-T2 half duplex mode
7	100BASE-T2 full duplex mode
8	PAUSE for full-duplex links
9	Asymmetric PAUSE for full-duplex links
10	Symmetric PAUSE for full-duplex links
11	Asymmetric and Symmetric PAUSE for full-duplex links
12	1000BASE-X, -LX, -SX, -CX half duplex mode
13	1000BASE-X, -LX, -SX, -CX full duplex mode
14	1000BASE-T half duplex mode
15	1000BASE-T full duplex mode

- ◆ **Remote Port Auto-Neg Status** – Shows whether port auto-negotiation is enabled on a port associated with the remote system.
- ◆ **Remote Port MAU Type** – An integer value that indicates the operational MAU type of the sending device. This object contains the integer value derived from the list position of the corresponding dot3MauType as listed in IETF RFC 3636 and is equal to the last number in the respective dot3MauType OID.

Port Details – 802.3 Extension Power Information

- ◆ **Remote Power Class** – The port Class of the given port associated with the remote system (PSE – Power Sourcing Equipment or PD – Powered Device).
- ◆ **Remote Power MDI Status** – Shows whether MDI power is enabled on the given port associated with the remote system.

- ◆ **Remote Power Pairs** – “Signal” means that the signal pairs only are in use, and “Spare” means that the spare pairs only are in use.
- ◆ **Remote Power MDI Supported** – Shows whether MDI power is supported on the given port associated with the remote system.
- ◆ **Remote Power Pair Controllable** – Indicates whether the pair selection can be controlled for sourcing power on the given port associated with the remote system.
- ◆ **Remote Power Classification** – This classification is used to tag different terminals on the Power over LAN network according to their power consumption. Devices such as IP telephones, WLAN access points and others, will be classified according to their power requirements.

Port Details – 802.3 Extension Trunk Information

- ◆ **Remote Link Aggregation Capable** – Shows if the remote port is not in link aggregation state and/or it does not support link aggregation.
- ◆ **Remote Link Aggregation Status** – The current aggregation status of the link.
- ◆ **Remote Link Port ID** – This object contains the IEEE 802.3 aggregated port identifier, aAggPortID (IEEE 802.3-2002, 30.7.2.1.1), derived from the ifNumber of the ifIndex for the port component associated with the remote system. If the remote port is not in link aggregation state and/or it does not support link aggregation, this value should be zero.

Port Details – 802.3 Extension Frame Information

- ◆ **Remote Max Frame Size** – An integer value indicating the maximum supported frame size in octets on the port component associated with the remote system.

Port Details – LLDP-MED Capability⁸

- ◆ **Device Class** – Any of the following categories of endpoint devices:
 - Class 1 – The most basic class of endpoint devices.
 - Class 2 – Endpoint devices that supports media stream capabilities.
 - Class 3 – Endpoint devices that directly supports end users of the IP communication systems.
 - Network Connectivity Device – Devices that provide access to the IEEE 802 based LAN infrastructure for LLDP-MED endpoint devices. These may be any LAN access device including LAN switch/router, IEEE 802.1 bridge, IEEE 802.3 repeater, IEEE 802.11 wireless access point, or any device that supports the IEEE 802.1AB and MED extensions defined by this Standard and can relay IEEE 802 frames via any method.

8. These fields are only displayed for end-node devices advertising LLDP-MED TLVs.

- ◆ **Supported Capabilities** – The supported set of capabilities that define the primary function(s) of the port:
 - LLDP-MED Capabilities
 - Network Policy
 - Location Identification
 - Extended Power via MDI – PSE
 - Extended Power via MDI – PD
 - Inventory
- ◆ **Current Capabilities** – The set of capabilities that define the primary function(s) of the port which are currently enabled.

*Port Details – Network Policy*⁸

- ◆ **Application Type** – The primary application(s) defined for this network policy:
 - Voice
 - Voice Signaling
 - Guest Signaling
 - Guest Voice Signaling
 - Softphone Voice
 - Video Conferencing
 - Streaming Video
 - Video Signaling
- ◆ **Tagged Flag** – Indicates whether the specified application type is using a tagged or untagged VLAN.
- ◆ **Layer 2 Priority** – The Layer 2 priority to be used for the specified application type. This field may specify one of eight priority levels (0-7), where a value of 0 represents use of the default priority.
- ◆ **Unknown Policy Flag** – Indicates that an endpoint device wants to explicitly advertise that this policy is required by the device, but is currently unknown.
- ◆ **VLAN ID** – The VLAN identifier (VID) for the port as defined in IEEE 802.1Q. A value of zero indicates that the port is using priority tagged frames, meaning that only the IEEE 802.1D priority level is significant and the default PVID of the ingress port is used instead.
- ◆ **DSCP Value** – The DSCP value to be used to provide Diffserv node behavior for the specified application type. This field may contain one of 64 code point values (0-63). A value of 0 represents use of the default DSCP value as defined in RFC 2475.

Port Details – Location Identification⁸

- ◆ **Location Data Format** – Any of these location ID data formats:
 - Coordinate-based LCI⁹ – Defined in RFC 3825, includes latitude resolution, latitude, longitude resolution, longitude, altitude type, altitude resolution, altitude, and datum.
 - Civic Address LCI⁹ – Includes What, Country code, CA type, CA length and CA value. “What” is described as the field entry “Device entry refers to” under “[Configuring LLDP Interface Attributes](#).” The other items and described under “[Configuring LLDP Interface Civic-Address](#).”
 - ECS ELIN – Emergency Call Service Emergency Location Identification Number supports traditional PSAP-based Emergency Call Service in North America.
- ◆ **Country Code** – The two-letter ISO 3166 country code in capital ASCII letters. (Example: DK, DE or US)
- ◆ **What** – The type of device to which the location applies as described for the field entry “Device entry refers to” under “[Configuring LLDP Interface Attributes](#).”

Port Details – Extended Power-via-MDI

- ◆ **Power Type** – Power Sourcing Entity (PSE) or Power Device (PD).
- ◆ **Power Priority** – Shows power priority for a port. (Unknown, Low, High, Critical)
- ◆ **Power Source** – Shows information based on the type of device:
 - **PD** – Unknown, PSE, Local, PSE and Local
 - **PSE** – Unknown, Primary Power Source, Backup Power Source - Power conservation mode
- ◆ **Power Value** – The total power in watts required by a PD device from a PSE device, or the total power a PSE device is capable of sourcing over a maximum length cable based on its current configuration. This parameter supports a maximum power required or available value of 102.3 Watts to allow for future expansion. (Range: 0 - 102.3 Watts)

Port Details – Inventory⁸

- ◆ **Hardware Revision** – The hardware revision of the end-point device.
- ◆ **Software Revision** – The software revision of the end-point device.

9. Location Configuration Information

- ◆ **Manufacture Name** – The manufacturer of the end-point device.
- ◆ **Asset ID** – The asset identifier of the end-point device. End-point devices are typically assigned asset identifiers to facilitate inventory management and assets tracking.
- ◆ **Firmware Revision** – The firmware revision of the end-point device.
- ◆ **Serial Number** – The serial number of the end-point device.
- ◆ **Model Name** – The model name of the end-point device.

Web Interface

To display LLDP information for a remote port:

1. Click Administration, LLDP.
2. Select Show Remote Device Information from the Step list.
3. Select Port, Port Details, Trunk, or Trunk Details.
4. When the next page opens, select a port on this switch and the index for a remote device attached to this port.
5. Click Query.

Figure 203: Displaying Remote Device Information for LLDP (Port)

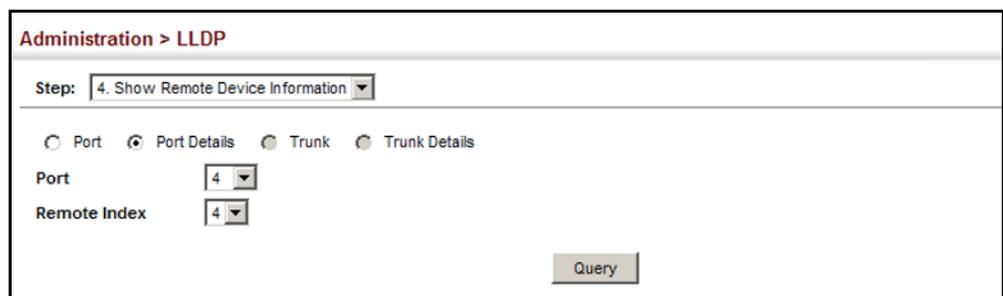
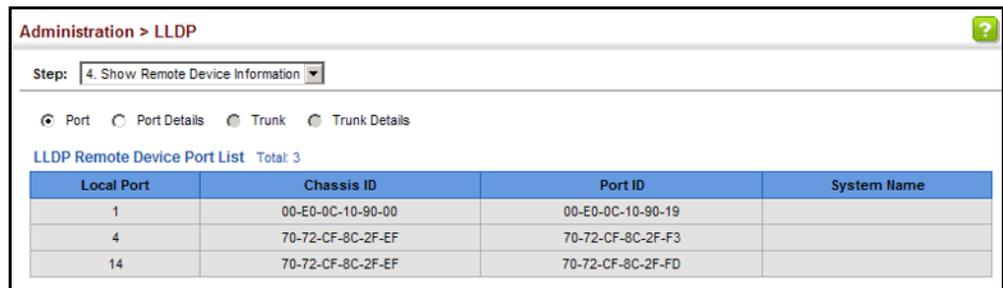


Figure 204: Displaying Remote Device Information for LLDP (Port Details)

Administration > LLDP

Step: 4. Show Remote Device Information

Port
 Port Details
 Trunk
 Trunk Details

Port: 54

Remote Index: 2

Query

LLDP Remote Device Port Information

Local Port	54	Port Type	MAC Address
Chassis Type	MAC Address	Port Description	Ethernet Port on unit 1, port 31
Chassis ID	70-72-CF-80-0E-50	Port ID	70-72-CF-80-0E-6F
System Name		System Capabilities Supported	Bridge, Router
System Description	AOS6700-32X	System Capabilities Enabled	Bridge, Router

Management Address List Total: 1

Address	Address Type
70-72-CF-80-0E-50	MAC Address

802.1 Extension Information

Remote Port VID: 1

Remote Port-Protocol VLAN List Total: 1

VLAN	Support	Status
3	Yes	Enabled

Remote VLAN Name List Total: 3

VLAN	Name
1	DefaultVlan
2	R&D
3	Protocol

Remote Protocol Identity List Total: 1

Remote Protocol Identity
88-CC

802.3 Extension Port Information

Remote Port Auto-Neg Supported	Yes	Remote Port Auto-Neg Status	Enabled
Remote Port Auto-Neg Adv-Capability (Hex)	0000	Remote Port MAU Type	35

802.3 Extension Trunk Information

Remote Link Aggregation Capable	Yes	Remote Link Aggregation Status	Disabled
Remote Link Port ID	0		

802.3 Extension Frame Information

Remote Max Frame Size	1522
------------------------------	------

Additional information displayed by an end-point device which advertises LLDP-MED TLVs is shown in the following figure.

Figure 205: Displaying Remote Device Information for LLDP (End Node)

Administration > LLDP			
Step: 4. Show Remote Device Information			
LLDP-MED Capability			
Device Class	Network Connectivity		
Supported Capabilities	LLDP-MED Capabilities, Network Policy, Location Identification, Inventory		
Current Capabilities	LLDP-MED Capabilities, Network Policy, Location Identification, Inventory		
Network Policy			
Application Type	Guest Voice Signaling	Unknown Policy Flag	Disabled
Tagged Flag	Disabled	VLAN ID	7
Layer 2 Priority	2	DSCP Value	62
Location Identification			
Location Data Format	Coordinate-based LCI		
Country Code	TW	What	2
Extended Power-via-MDI			
Power Type	PSE	Power Source	Unknown
Power Priority	Unknown	Power Value	0 W Watts
Inventory			
Hardware Revision	R01	Firmware Revision	1.0.0.2
Software Revision	1.0.0.2	Serial Number	
Manufacture Name		Model Name	
Asset ID	1		

Displaying Device Statistics Use the Administration > LLDP (Show Device Statistics) page to display statistics for LLDP-capable devices attached to the switch, and for LLDP protocol messages transmitted or received on all local interfaces.

Parameters

These parameters are displayed:

General Statistics on Remote Devices

- ◆ **Neighbor Entries List Last Updated** – The time the LLDP neighbor entry list was last updated.
- ◆ **New Neighbor Entries Count** – The number of LLDP neighbors for which the remote TTL has not yet expired.
- ◆ **Neighbor Entries Deleted Count** – The number of LLDP neighbors which have been removed from the LLDP remote systems MIB for any reason.
- ◆ **Neighbor Entries Dropped Count** – The number of times which the remote database on this switch dropped an LLDPDU because of insufficient resources.

- ◆ **Neighbor Entries Age-out Count** – The number of times that a neighbor’s information has been deleted from the LLDP remote systems MIB because the remote TTL timer has expired.

Port/Trunk

- ◆ **Frames Discarded** – Number of frames discarded because they did not conform to the general validation rules as well as any specific usage rules defined for the particular TLV.
- ◆ **Frames Invalid** – A count of all LLDPDUs received with one or more detectable errors.
- ◆ **Frames Received** – Number of LLDP PDUs received.
- ◆ **Frames Sent** – Number of LLDP PDUs transmitted.
- ◆ **TLVs Unrecognized** – A count of all TLVs not recognized by the receiving LLDP local agent.
- ◆ **TLVs Discarded** – A count of all LLDPDUs received and then discarded due to insufficient memory space, missing or out-of-sequence attributes, or any other reason.
- ◆ **Neighbor Ageouts** – A count of the times that a neighbor’s information has been deleted from the LLDP remote systems MIB because the remote TTL timer has expired.

Web Interface

To display statistics for LLDP-capable devices attached to the switch:

1. Click Administration, LLDP.
2. Select Show Device Statistics from the Step list.
3. Select General, Port, or Trunk.

Figure 206: Displaying LLDP Device Statistics (General)

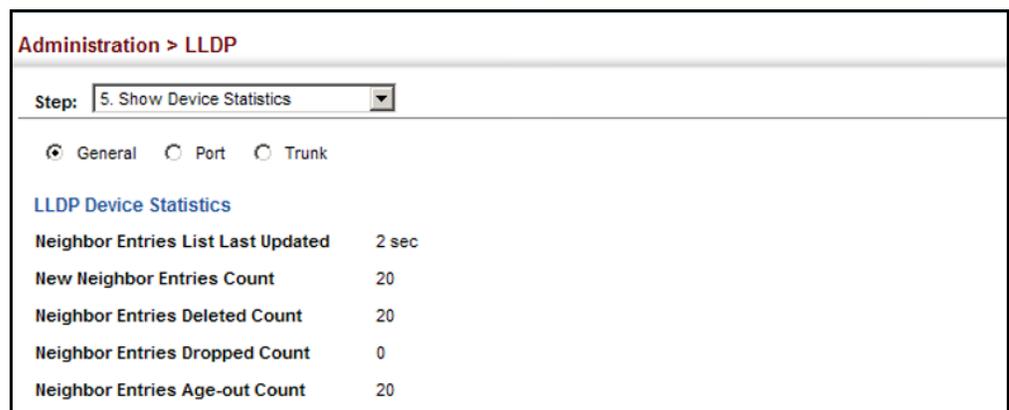
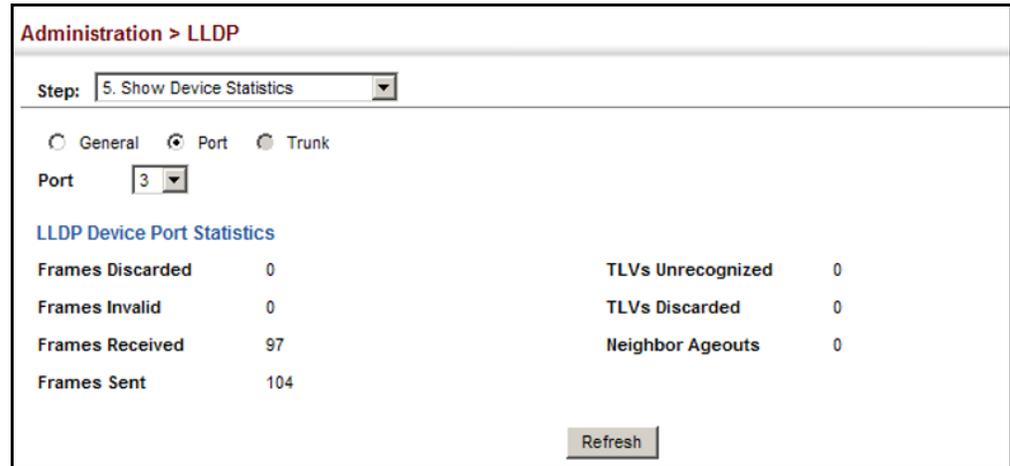


Figure 207: Displaying LLDP Device Statistics (Port)



Simple Network Management Protocol

Simple Network Management Protocol (SNMP) is a communication protocol designed specifically for managing devices on a network. Equipment commonly managed with SNMP includes switches, routers and host computers. SNMP is typically used to configure these devices for proper operation in a network environment, as well as to monitor them to evaluate performance or detect potential problems.

Managed devices supporting SNMP contain software, which runs locally on the device and is referred to as an agent. A defined set of variables, known as managed objects, is maintained by the SNMP agent and used to manage the device. These objects are defined in a Management Information Base (MIB) that provides a standard presentation of the information controlled by the agent. SNMP defines both the format of the MIB specifications and the protocol used to access this information over the network.

The switch includes an onboard agent that supports SNMP versions 1, 2c, and 3. This agent continuously monitors the status of the switch hardware, as well as the traffic passing through its ports. A network management station can access this information using network management software. Access to the onboard agent from clients using SNMP v1 and v2c is controlled by community strings. To communicate with the switch, the management station must first submit a valid community string for authentication.

Access to the switch from clients using SNMPv3 provides additional security features that cover message integrity, authentication, and encryption; as well as controlling user access to specific areas of the MIB tree.

The SNMPv3 security structure consists of security models, with each model having its own security levels. There are three security models defined, SNMPv1, SNMPv2c, and SNMPv3. Users are assigned to “groups” that are defined by a security model

and specified security levels. Each group also has a defined security access to set of MIB objects for reading and writing, which are known as “views.” The switch has a default view (all MIB objects) and default groups defined for security models v1 and v2c. The following table shows the security models and levels available and the system default settings.

Table 32: SNMPv3 Security Models and Levels

Model	Level	Group	Read View	Write View	Notify View	Security
v1	noAuthNoPriv	public (read only)	defaultview	none	none	Community string only
v1	noAuthNoPriv	private (read/write)	defaultview	defaultview	none	Community string only
v1	noAuthNoPriv	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	Community string only
v2c	noAuthNoPriv	public (read only)	defaultview	none	none	Community string only
v2c	noAuthNoPriv	private (read/write)	defaultview	defaultview	none	Community string only
v2c	noAuthNoPriv	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	Community string only
v3	noAuthNoPriv	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	A user name match only
v3	AuthNoPriv	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	Provides user authentication via MD5 or SHA algorithms
v3	AuthPriv	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	Provides user authentication via MD5 or SHA algorithms and data privacy using DES 56-bit encryption



Note: The predefined default groups and view can be deleted from the system. You can then define customized groups and views for the SNMP clients that require access.

Command Usage

Configuring SNMPv1/2c Management Access

To configure SNMPv1 or v2c management access to the switch, follow these steps:

1. Use the Administration > SNMP (Configure Global) page to enable SNMP on the switch, and to enable trap messages.
2. Use the Administration > SNMP (Configure User - Add Community) page to configure the community strings authorized for management access.
3. Use the Administration > SNMP (Configure Trap) page to specify trap managers so that key events are reported by this switch to your management station.

Configuring SNMPv3 Management Access

1. Use the Administration > SNMP (Configure Global) page to enable SNMP on the switch, and to enable trap messages.
2. Use the Administration > SNMP (Configure Trap) page to specify trap managers so that key events are reported by this switch to your management station.
3. Use the Administration > SNMP (Configure Engine) page to change the local engine ID. If you want to change the default engine ID, it must be changed before configuring other parameters.
4. Use the Administration > SNMP (Configure View) page to specify read and write access views for the switch MIB tree.
5. Use the Administration > SNMP (Configure User) page to configure SNMP user groups with the required security model (i.e., SNMP v1, v2c or v3) and security level (i.e., authentication and privacy).
6. Use the Administration > SNMP (Configure Group) page to assign SNMP users to groups, along with their specific authentication and privacy passwords.

Configuring Global Settings for SNMP

Use the Administration > SNMP (Configure Global) page to enable SNMPv3 service for all management clients (i.e., versions 1, 2c, 3), and to enable trap messages.

Parameters

These parameters are displayed:

- ◆ **Agent Status** – Enables SNMP on the switch. (Default: Enabled)
- ◆ **Authentication Traps**¹⁰ – Issues a notification message to specified IP trap managers whenever an invalid community string is submitted during the SNMP access authentication process. (Default: Enabled)

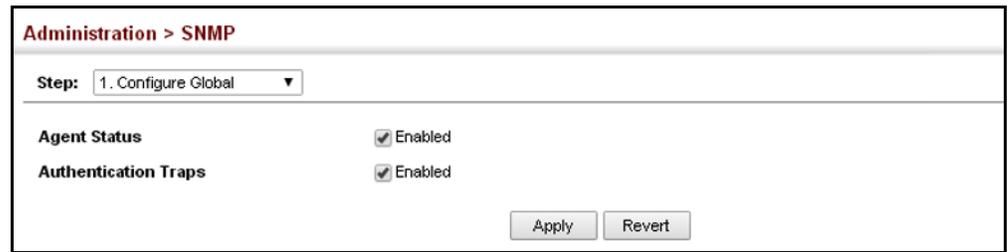
Web Interface

To configure global settings for SNMP:

1. Click Administration, SNMP.
2. Select Configure Global from the Step list.
3. Enable SNMP and the required trap types.
4. Click Apply

10. These are legacy notifications and therefore when used for SNMPv3 hosts, they must be enabled in conjunction with the corresponding entries in the Notification View ([page 358](#)).

Figure 208: Configuring Global Settings for SNMP



Setting the Local Engine ID

Use the Administration > SNMP (Configure Engine - Set Engine ID) page to change the local engine ID. An SNMPv3 engine is an independent SNMP agent that resides on the switch. This engine protects against message replay, delay, and redirection. The engine ID is also used in combination with user passwords to generate the security keys for authenticating and encrypting SNMPv3 packets.

Command Usage

- ◆ A local engine ID is automatically generated that is unique to the switch. This is referred to as the default engine ID. If the local engine ID is deleted or changed, all SNMP users will be cleared. You will need to reconfigure all existing users.

Parameters

These parameters are displayed:

- ◆ **Engine ID** – A new engine ID can be specified by entering 9 to 64 hexadecimal characters (5 to 32 octets in hexadecimal format). If an odd number of characters are specified, a trailing zero is added to the value to fill in the last octet. For example, the value “123456789” is equivalent to “1234567890”.
- ◆ **Engine Boots** – The number of times that the engine has (re-)initialized since the SNMP Engine ID was last configured.

Web Interface

To configure the local SNMP engine ID:

1. Click Administration, SNMP.
2. Select Configure Engine from the Step list.
3. Select Set Engine ID from the Action list.
4. Enter an ID of a least 9 hexadecimal characters.
5. Click Apply

Figure 209: Configuring the Local Engine ID for SNMP

The screenshot shows a web interface for configuring SNMP. At the top, it says 'Administration > SNMP'. Below that, there are two dropdown menus: 'Step: 2. Configure Engine' and 'Action: Set Engine ID'. The main configuration area has two fields: 'Engine ID' with the value '800001030300000c0000fd0000' and 'Engine Boots' with the value '5'. At the bottom right, there are two buttons: 'Default' and 'Save'.

Specifying a Remote Engine ID

Use the Administration > SNMP (Configure Engine - Add Remote Engine) page to configure a engine ID for a remote management station. To allow management access from an SNMPv3 user on a remote device, you must first specify the engine identifier for the SNMP agent on the remote device where the user resides. The remote engine ID is used to compute the security digest for authentication and encryption of packets passed between the switch and a user on the remote host.

Command Usage

- ◆ SNMP passwords are localized using the engine ID of the authoritative agent. For informs, the authoritative SNMP agent is the remote agent. You therefore need to configure the remote agent's SNMP engine ID before you can send proxy requests or informs to it. (See ["Configuring Remote SNMPv3 Users" on page 369.](#))

Parameters

These parameters are displayed:

- ◆ **Remote Engine ID** – The engine ID can be specified by entering 9 to 64 hexadecimal characters (5 to 32 octets in hexadecimal format). If an odd number of characters are specified, a trailing zero is added to the value to fill in the last octet. For example, the value "123456789" is equivalent to "1234567890".
- ◆ **Remote IP Host** – The IPv4 address of a remote management station which is using the specified engine ID.

Web Interface

To configure a remote SNMP engine ID:

1. Click Administration, SNMP.
2. Select Configure Engine from the Step list.
3. Select Add Remote Engine from the Action list.
4. Enter an ID of a least 9 hexadecimal characters, and the IP address of the remote host.
5. Click Apply

Figure 210: Configuring a Remote Engine ID for SNMP

Administration > SNMP

Step: 2. Configure Engine Action: Add Remote Engine

Remote Engine ID: 5432100000

Remote IP Host: 192.168.1.19

Apply Revert

To show the remote SNMP engine IDs:

1. Click Administration, SNMP.
2. Select Configure Engine from the Step list.
3. Select Show Remote Engine from the Action list.

Figure 211: Showing Remote Engine IDs for SNMP

Administration > SNMP

Step: 2. Configure Engine Action: Show Remote Engine

SNMPv3 Remote Engine List Total: 2

<input type="checkbox"/>	Remote Engine ID	Remote IP Host
<input type="checkbox"/>	1234567890	1.2.3.4
<input type="checkbox"/>	0a9b8c7d6e5f	5.6.7.8

Delete Revert

Setting SNMPv3 Views Use the Administration > SNMP (Configure View) page to configure SNMPv3 views which are used to restrict user access to specified portions of the MIB tree. The predefined view “defaultview” includes access to the entire MIB tree.

Parameters

These parameters are displayed:

Add View

- ◆ **View Name** – Lists the SNMP views configured in the Add View page. A maximum of 32 views can be configured. (Range: 1-32 characters)
- ◆ **OID Subtree** – Specifies the initial object identifier of a branch within the MIB tree. Wild cards can be used to mask a specific portion of the OID string. Use the Add OID Subtree page to configure additional object identifiers. (Range: 1-64 characters)
- ◆ **Type** – Indicates if the object identifier of a branch within the MIB tree is included or excluded from the SNMP view.

Add OID Subtree

- ◆ **View Name** – Lists the SNMP views configured in the Add View page. A maximum of 32 views can be configured. (Range: 1-32 characters)
- ◆ **OID Subtree** – Adds an additional object identifier of a branch within the MIB tree to the selected View. Wild cards can be used to mask a specific portion of the OID string. (Range: 1-64 characters)
- ◆ **Type** – Indicates if the object identifier of a branch within the MIB tree is included or excluded from the SNMP view.

Web Interface

To configure an SNMP view of the switch's MIB database:

1. Click Administration, SNMP.
2. Select Configure View from the Step list.
3. Select Add View from the Action list.
4. Enter a view name and specify the initial OID subtree in the switch's MIB database to be included or excluded in the view. Use the Add OID Subtree page to add additional object identifier branches to the view.
5. Click Apply

Figure 212: Creating an SNMP View

Administration > SNMP

Step: 3. Configure View Action: Add View

View Name: ifEntry.a

OID Subtree: 1.3.6.1.2.1.2.2.1.1.*

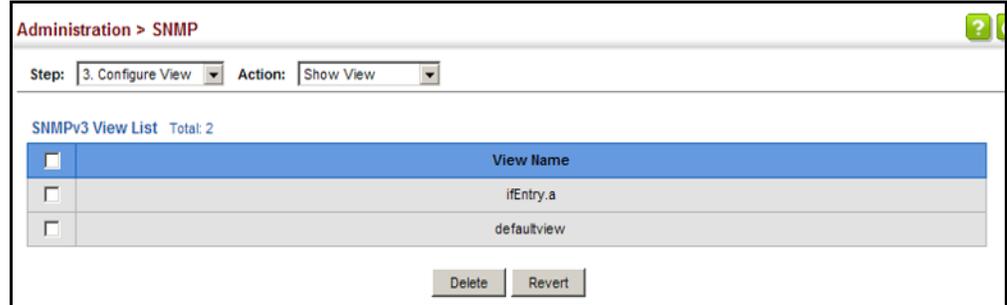
Type: Included

Apply Revert

To show the SNMP views of the switch's MIB database:

1. Click Administration, SNMP.
2. Select Configure View from the Step list.
3. Select Show View from the Action list.

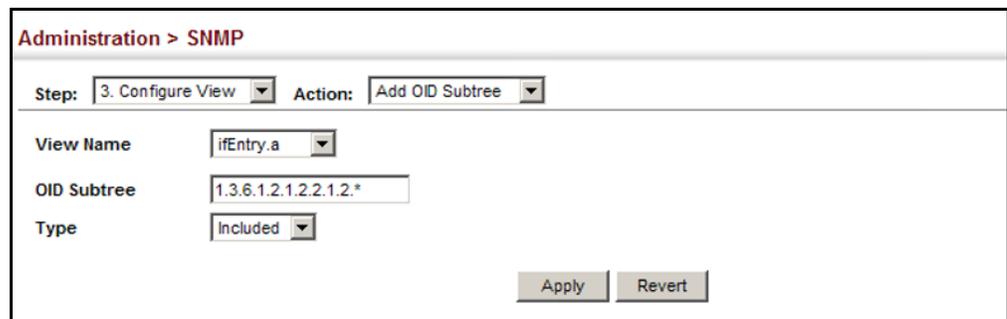
Figure 213: Showing SNMP Views



To add an object identifier to an existing SNMP view of the switch's MIB database:

1. Click Administration, SNMP.
2. Select Configure View from the Step list.
3. Select Add OID Subtree from the Action list.
4. Select a view name from the list of existing views, and specify an additional OID subtree in the switch's MIB database to be included or excluded in the view.
5. Click Apply

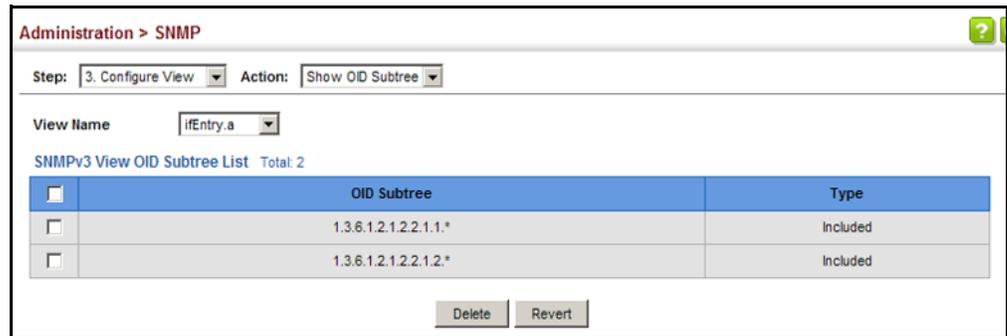
Figure 214: Adding an OID Subtree to an SNMP View



To show the OID branches configured for the SNMP views of the switch's MIB database:

1. Click Administration, SNMP.
2. Select Configure View from the Step list.
3. Select Show OID Subtree from the Action list.
4. Select a view name from the list of existing views.

Figure 215: Showing the OID Subtree Configured for SNMP Views



Configuring SNMPv3 Groups

Use the Administration > SNMP (Configure Group) page to add an SNMPv3 group which can be used to set the access policy for its assigned users, restricting them to specific read, write, and notify views. You can use the pre-defined default groups or create new groups to map a set of SNMP users to SNMP views.

Parameters

These parameters are displayed:

- ◆ **Group Name** – The name of the SNMP group to which the user is assigned. A maximum of 22 groups can be configured. (Range: 1-32 characters)
- ◆ **Security Model** – The user security model; SNMP v1, v2c or v3.
- ◆ **Security Level** – The following security levels are only used for the groups assigned to the SNMP security model:
 - **noAuthNoPriv** – There is no authentication or encryption used in SNMP communications. (This is the default security level.)
 - **AuthNoPriv** – SNMP communications use authentication, but the data is not encrypted.
 - **AuthPriv** – SNMP communications use both authentication and encryption.
- ◆ **Read View** – The configured view for read access. (Range: 1-32 characters)
- ◆ **Write View** – The configured view for write access. (Range: 1-32 characters)
- ◆ **Notify View** – The configured view for notifications. (Range: 1-32 characters)

Table 33: Supported Notification Messages

Model	Level	Group
<i>RFC 1493 Traps</i>		
newRoot	1.3.6.1.2.1.17.0.1	The newRoot trap indicates that the sending agent has become the new root of the Spanning Tree; the trap is sent by a bridge soon after its election as the new root, e.g., upon expiration of the Topology Change Timer immediately subsequent to its election.
topologyChange	1.3.6.1.2.1.17.0.2	A topologyChange trap is sent by a bridge when any of its configured ports transitions from the Learning state to the Forwarding state, or from the Forwarding state to the Discarding state. The trap is not sent if a newRoot trap is sent for the same transition.
<i>SNMPv2 Traps</i>		
coldStart	1.3.6.1.6.3.1.1.5.1	A coldStart trap signifies that the SNMPv2 entity, acting in an agent role, is reinitializing itself and that its configuration may have been altered.
warmStart	1.3.6.1.6.3.1.1.5.2	A warmStart trap signifies that the SNMPv2 entity, acting in an agent role, is reinitializing itself such that its configuration is unaltered.
linkDown*	1.3.6.1.6.3.1.1.5.3	A linkDown trap signifies that the SNMP entity, acting in an agent role, has detected that the ifOperStatus object for one of its communication links is about to enter the down state from some other state (but not from the notPresent state). This other state is indicated by the included value of ifOperStatus.
linkUp*	1.3.6.1.6.3.1.1.5.4	A linkUp trap signifies that the SNMP entity, acting in an agent role, has detected that the ifOperStatus object for one of its communication links left the down state and transitioned into some other state (but not into the notPresent state). This other state is indicated by the included value of ifOperStatus.
authenticationFailure*	1.3.6.1.6.3.1.1.5.5	An authenticationFailure trap signifies that the SNMPv2 entity, acting in an agent role, has received a protocol message that is not properly authenticated. While all implementations of the SNMPv2 must be capable of generating this trap, the snmpEnableAuthenTraps object indicates whether this trap will be generated.
<i>RMON Events (V2)</i>		
risingAlarm	1.3.6.1.2.1.16.0.1	The SNMP trap that is generated when an alarm entry crosses its rising threshold and generates an event that is configured for sending SNMP traps.
fallingAlarm	1.3.6.1.2.1.16.0.2	The SNMP trap that is generated when an alarm entry crosses its falling threshold and generates an event that is configured for sending SNMP traps.
<i>Private Traps</i>		
swPowerStatus ChangeTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.1	This trap is sent when the power state changes.
swFanFailureTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.17	This trap is sent when the fan fails.
swFanRecoverTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.18	This trap is sent when fan failure has recovered.
swThermalRisingNotification	1.3.6.1.4.1.259.12.1.2.2.1.0.58	This trap is sent when the temperature is over the switchThermalActionRisingThreshold.

Table 33: Supported Notification Messages (Continued)

Model	Level	Group
swThermalFallingNotification	1.3.6.1.4.1.259.12.1.2.2.1.0.59	This trap is sent when the temperature is below the switchThermalActionFallingThreshold.
dot1agCfmMepUpTrap	1.3.6.1.4.1.259.10.1.36.2.1.0.97	This trap is sent when a new remote MEP is discovered.
dot1agCfmMepDownTrap	1.3.6.1.4.1.259.10.1.36.2.1.0.98	This trap is sent when port status or interface status TLV received from a remote MEP indicates it is not up.
dot1agCfmConfigFailTrap	1.3.6.1.4.1.259.10.1.36.2.1.0.99	This trap is sent when a MEP receives a CCM with an MPID which already exists on the same MA in this switch.
dot1agCfmLoopFindTrap	1.3.6.1.4.1.259.10.1.36.2.1.0.100	This trap is sent when a MEP receives its own CCMs.
dot1agCfmMepUnknownTrap	1.3.6.1.4.1.259.10.1.36.2.1.0.101	This trap is sent when a CCM is received from an unexpected MEP.
dot1agCfmMepMissingTrap	1.3.6.1.4.1.259.10.1.36.2.1.0.102	This trap is sent when the cross-check enable timer expires and no CCMs were received from an expected (configured) MEP.
dot1agCfmMaUpTrap	1.3.6.1.4.1.259.10.1.36.2.1.0.103	This trap is sent when all expected remote MEPs are up.
autoUpgradeTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.104	This trap is sent when auto upgrade is executed.
swCpuUtiRisingNotification	1.3.6.1.4.1.259.12.1.2.2.1.0.107	This notification indicates that the CPU utilization has risen from cpuUtiFallingThreshold to cpuUtiRisingThreshold.
swCpuUtiFallingNotification	1.3.6.1.4.1.259.12.1.2.2.1.0.108	This notification indicates that the CPU utilization has fallen from cpuUtiRisingThreshold to cpuUtiFallingThreshold.
swMemoryUtiRisingThreshold Notification	1.3.6.1.4.1.259.12.1.2.2.1.0.109	This notification indicates that the memory utilization has risen from memoryUtiFallingThreshold to memoryUtiRisingThreshold.
swMemoryUtiFallingThreshold Notification	1.3.6.1.4.1.259.12.1.2.2.1.0.110	This notification indicates that the memory utilization has fallen from memoryUtiRisingThreshold to memoryUtiFallingThreshold.
dhcpRogueServerAttackTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.114	This trap is sent when receiving a DHCP packet from a rouge server.
macNotificationTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.138	This trap is sent when there are changes of the dynamic MAC addresses on the switch.
sfpThresholdAlarmWarnTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.189	This trap is sent when the SFP's A/D quantity is not within alarm/warning thresholds.
udldPortShutdownTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.192	This trap is sent when the port is shut down by UDLD.
userAuthenticationFailureTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.199	This trap will be triggered if authentication is fail.
userAuthenticationSuccessTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.200	This trap will be triggered if authentication is successful.
loginTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.201	This trap is sent when user login.

Table 33: Supported Notification Messages (Continued)

Model	Level	Group
logoutTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.202	This trap is sent when user logout.
fileCopyTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.208	This trap is sent when file copy is executed.
userauthCreateUserTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.209	This trap is sent when create user account.
userauthDeleteUserTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.210	This trap is sent when delete user account.
userauthModifyUserPrivilegeTrap	1.3.6.1.4.1.259.12.1.2.2.1.0.211	This trap is sent when modify user privilege.

* These are legacy notifications and therefore must be enabled in conjunction with the corresponding traps on the SNMP Configuration menu.

Web Interface

To configure an SNMP group:

1. Click Administration, SNMP.
2. Select Configure Group from the Step list.
3. Select Add from the Action list.
4. Enter a group name, assign a security model and level, and then select read, write, and notify views.
5. Click Apply

Figure 216: Creating an SNMP Group

To show SNMP groups:

1. Click Administration, SNMP.
2. Select Configure Group from the Step list.
3. Select Show from the Action list.

Figure 217: Showing SNMP Groups

<input type="checkbox"/>	Group Name	Model	Level	Read View	Write View	Notify View
<input type="checkbox"/>	public	v1	noAuthNoPriv	defaultview	No writeview specified	No notifyview specified
<input type="checkbox"/>	public	v2c	noAuthNoPriv	defaultview	No writeview specified	No notifyview specified
<input type="checkbox"/>	private	v1	noAuthNoPriv	defaultview	defaultview	No notifyview specified
<input type="checkbox"/>	private	v2c	noAuthNoPriv	defaultview	defaultview	No notifyview specified
<input type="checkbox"/>	secure-users	v3	authNoPriv	ifEntry.a	ifEntry.a	ifEntry.a

Setting Community Access Strings Use the Administration > SNMP (Configure User - Add Community) page to configure up to five community strings authorized for management access by clients using SNMP v1 and v2c. For security reasons, you should consider removing the default strings.

Parameters

These parameters are displayed:

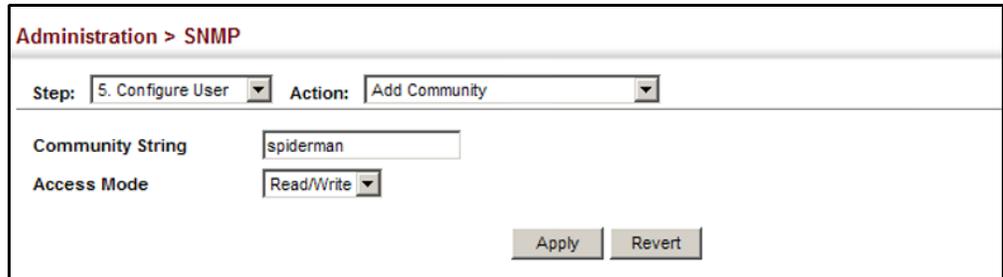
- ◆ **Community String** – A community string that acts like a password and permits access to the SNMP protocol.
Range: 1-32 characters, case sensitive
Default strings: “public” (Read-Only), “private” (Read/Write)
- ◆ **Access Mode** – Specifies the access rights for the community string:
 - **Read-Only** – Authorized management stations are only able to retrieve MIB objects.
 - **Read/Write** – Authorized management stations are able to both retrieve and modify MIB objects.

Web Interface

To set a community access string:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.
3. Select Add Community from the Action list.
4. Add new community strings as required, and select the corresponding access rights from the Access Mode list.
5. Click Apply

Figure 218: Setting Community Access Strings

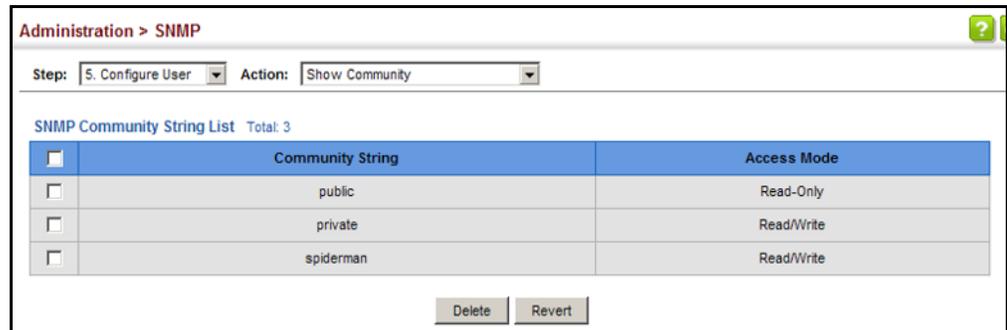


To show the community access strings:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.

3. Select Show Community from the Action list.

Figure 219: Showing Community Access Strings



Configuring Local SNMPv3 Users

Use the Administration > SNMP (Configure User - Add SNMPv3 Local User) page to authorize management access for SNMPv3 clients, or to identify the source of SNMPv3 trap messages sent from the local switch. Each SNMPv3 user is defined by a unique name. Users must be configured with a specific security level and assigned to a group. The SNMPv3 group restricts users to a specific read, write, and notify view.

Parameters

These parameters are displayed:

- ◆ **User Name** – The name of user connecting to the SNMP agent. A maximum of three local users can be configured. (Range: 1-32 characters)
- ◆ **Group Name** – The name of the SNMP group to which the user is assigned. (Range: 1-32 characters)
- ◆ **Security Model** – The user security model; SNMP v1, v2c or v3.
- ◆ **Security Level** – The following security levels are only used for the groups assigned to the SNMP security model:
 - **noAuthNoPriv** – There is no authentication or encryption used in SNMP communications. (This is the default security level.)
 - **AuthNoPriv** – SNMP communications use authentication, but the data is not encrypted.
 - **AuthPriv** – SNMP communications use both authentication and encryption.
- ◆ **Authentication Protocol** – The method used for user authentication. (Options: MD5, SHA; Default: MD5)
- ◆ **Authentication Password** – A minimum of eight plain text characters is required.

- ◆ **Privacy Protocol** – The encryption algorithm used for data privacy:
 - **3DES** - Uses SNMPv3 with privacy with 3DES (168-bit) encryption.
 - **AES128** - Uses SNMPv3 with privacy with AES128 encryption.
 - **AES192** - Uses SNMPv3 with privacy with AES192 encryption.
 - **AES256** - Uses SNMPv3 with privacy with AES256 encryption.
 - **DES56** - Uses SNMPv3 with privacy with DES56 encryption.

- ◆ **Privacy Password** – A minimum of eight plain text characters is required.

Web Interface

To configure a local SNMPv3 user:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.
3. Select Add SNMPv3 Local User from the Action list.
4. Enter a name and assign it to a group. If the security model is set to SNMPv3 and the security level is authNoPriv or authPriv, then an authentication protocol and password must be specified. If the security level is authPriv, a privacy password must also be specified.
5. Click Apply

Figure 220: Configuring Local SNMPv3 Users

Administration > SNMP

Step: 5. Configure User Action: Add SNMPv3 Local User

SNMPv3 User

User Name: chris

Group Name: public r&d

Security Model: v3

Security Level: authPriv

User Authentication

Authentication Protocol: MD5

Authentication Password: greenpeace

Data Privacy

Privacy Protocol: DES56

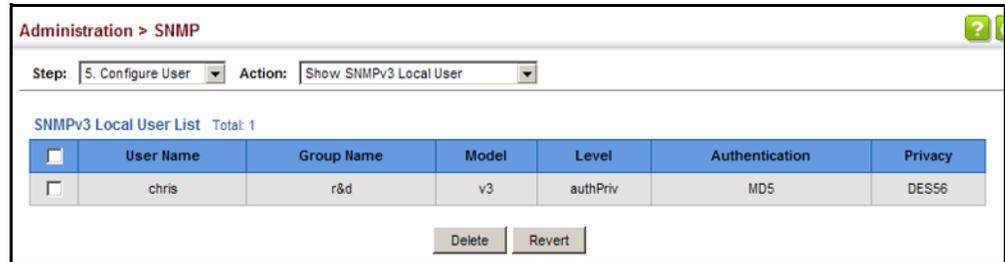
Privacy Password: einstien

Apply Revert

To show local SNMPv3 users:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.
3. Select Show SNMPv3 Local User from the Action list.

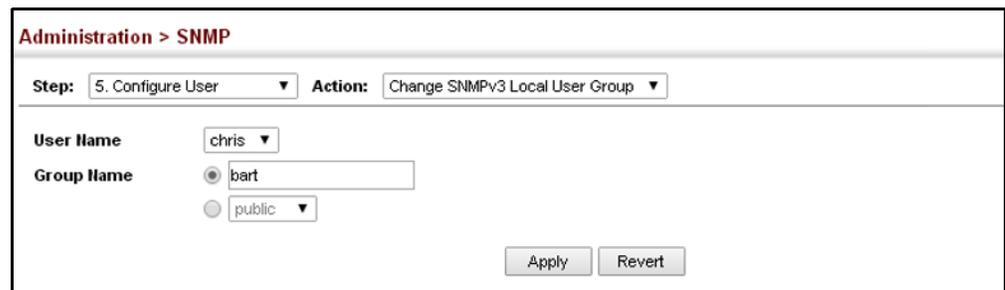
Figure 221: Showing Local SNMPv3 Users



To change a local SNMPv3 local user group:

1. Click Administration, SNMP.
2. Select Change SNMPv3 Local User Group from the Action list.
3. Select the User Name.
4. Enter a new group name.
5. Click Apply

Figure 222: Changing a Local SNMPv3 User Group



Configuring Remote SNMPv3 Users

Use the Administration > SNMP (Configure User - Add SNMPv3 Remote User) page to identify the source of SNMPv3 inform messages sent from the local switch. Each SNMPv3 user is defined by a unique name. Users must be configured with a specific security level and assigned to a group. The SNMPv3 group restricts users to a specific read, write, and notify view.

Command Usage

- ◆ To grant management access to an SNMPv3 user on a remote device, you must first specify the engine identifier for the SNMP agent on the remote device where the user resides. The remote engine ID is used to compute the security digest for authentication and encryption of packets passed between the switch and the remote user. (See [“Specifying Trap Managers” on page 372](#) and [“Specifying a Remote Engine ID” on page 357.](#))

Parameters

These parameters are displayed:

- ◆ **User Name** – The name of user connecting to the SNMP agent.
(Range: 1-32 characters)
- ◆ **Group Name** – The name of the SNMP group to which the user is assigned.
(Range: 1-32 characters)
- ◆ **Remote IP** – The IPv4 address of the remote device where the user resides.
- ◆ **Security Model** – The user security model; SNMP v1, v2c or v3. (Default: v3)
- ◆ **Security Level** – The following security levels are only used for the groups assigned to the SNMP security model:
 - **noAuthNoPriv** – There is no authentication or encryption used in SNMP communications. (This is the default security level.)
 - **AuthNoPriv** – SNMP communications use authentication, but the data is not encrypted.
 - **AuthPriv** – SNMP communications use both authentication and encryption.
- ◆ **Authentication Protocol** – The method used for user authentication.
(Options: MD5, SHA; Default: MD5)
- ◆ **Authentication Password** – A minimum of eight plain text characters is required.
- ◆ **Privacy Protocol** – The encryption algorithm used for data privacy:
 - **3DES** - Uses SNMPv3 with privacy with 3DES (168-bit) encryption.
 - **AES128** - Uses SNMPv3 with privacy with AES128 encryption.
 - **AES192** - Uses SNMPv3 with privacy with AES192 encryption.
 - **AES256** - Uses SNMPv3 with privacy with AES256 encryption.
 - **DES56** - Uses SNMPv3 with privacy with DES56 encryption.
- ◆ **Privacy Password** – A minimum of eight plain text characters is required.

Web Interface

To configure a remote SNMPv3 user:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.
3. Select Add SNMPv3 Remote User from the Action list.
4. Enter a name and assign it to a group. Enter the IP address to identify the source of SNMPv3 inform messages sent from the local switch. If the security model is set to SNMPv3 and the security level is authNoPriv or authPriv, then an authentication protocol and password must be specified. If the security level is authPriv, a privacy password must also be specified.
5. Click Apply

Figure 223: Configuring Remote SNMPv3 Users

The screenshot shows the 'Administration > SNMP' configuration page. At the top, there are two dropdown menus: 'Step: 5. Configure User' and 'Action: Add SNMPv3 Remote User'. Below this, the configuration is organized into sections:

- SNMPv3 User**
 - User Name: mark
 - Group Name: public (selected in dropdown), r&d (in text box)
 - Remote IP: 192.168.1.19
 - Security Model: v3
 - Security Level: authPriv
- User Authentication**
 - Authentication Protocol: MD5
 - Authentication Password: greenpeace
- Data Privacy**
 - Privacy Protocol: DES56
 - Privacy Password: einstien

At the bottom right, there are 'Apply' and 'Revert' buttons.

To show remote SNMPv3 users:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.
3. Select Show SNMPv3 Remote User from the Action list.

Figure 224: Showing Remote SNMPv3 Users

<input type="checkbox"/>	User Name	Group Name	Engine ID	Model	Level	Authentication	Privacy
<input type="checkbox"/>	mark	r&d	0123456789	v3	authPriv	MD5	DES56

Specifying Trap Managers

Use the Administration > SNMP (Configure Trap) page to specify the host devices to be sent traps and the types of traps to send. Traps indicating status changes are issued by the switch to the specified trap managers. You must specify trap managers so that key events are reported by this switch to your management station (using network management software). You can specify up to five management stations that will receive authentication failure messages and other trap messages from the switch.

Command Usage

- ◆ Notifications are issued by the switch as trap messages by default. The recipient of a trap message does not send a response to the switch. Traps are therefore not as reliable as inform messages, which include a request for acknowledgement of receipt. Informs can be used to ensure that critical information is received by the host. However, note that informs consume more system resources because they must be kept in memory until a response is received. Informs also add to network traffic. You should consider these effects when deciding whether to issue notifications as traps or informs.

To send an inform to a SNMPv2c host, complete these steps:

1. Enable the SNMP agent ([page 355](#)).
2. Create a view with the required notification messages ([page 358](#)).
3. Configure the group (matching the community string specified on the Configure Trap - Add page) to include the required notify view ([page 361](#)).
4. Enable trap informs as described in the following pages.

To send an inform to a SNMPv3 host, complete these steps:

1. Enable the SNMP agent ([page 355](#)).
2. Create a remote SNMPv3 user to use in the message exchange process ([page 367](#)). If the user specified in the trap configuration page does not exist, an SNMPv3 group will be automatically created using the name of the specified remote user, and default settings for the read, write, and notify view.
3. Create a view with the required notification messages ([page 358](#)).
4. Create a group that includes the required notify view ([page 361](#)).
5. Enable trap informs as described in the following pages.

Parameters

These parameters are displayed:

SNMP Version 1

- ◆ **IP Address** – IPv4 or IPv6 address of a new management station to receive notification message (i.e., the targeted recipient).
- ◆ **Version** – Specifies whether to send notifications as SNMP v1, v2c, or v3 traps. (Default: v1)
- ◆ **Community String** – Specifies a valid community string for the new trap manager entry. (Range: 1-32 characters, case sensitive)

Although you can set this string in the Configure Trap – Add page, we recommend defining it in the Configure User – Add Community page.

- ◆ **UDP Port** – Specifies the UDP port number used by the trap manager. (Default: 162)

SNMP Version 2c

- ◆ **IP Address** – IPv4 or IPv6 address of a new management station to receive notification message (i.e., the targeted recipient).
- ◆ **Version** – Specifies whether to send notifications using SNMP v1, v2c, or v3 traps.
- ◆ **Notification Type**
 - **Traps** – Notifications are sent as trap messages.
 - **Inform** – Notifications are sent as inform messages. Note that this option is only available for version 2c and 3 hosts. (Default: traps are used)
 - **Timeout** – The number of seconds to wait for an acknowledgment before resending an inform message. (Range: 0-2147483647 centiseconds; Default: 1500 centiseconds)
 - **Retry times** – The maximum number of times to resend an inform message if the recipient does not acknowledge receipt. (Range: 0-255; Default: 3)
- ◆ **Community String** – Specifies a valid community string for the new trap manager entry. (Range: 1-32 characters, case sensitive)

Although you can set this string in the Configure Trap – Add page, we recommend defining it in the Configure User – Add Community page.

- ◆ **UDP Port** – Specifies the UDP port number used by the trap manager. (Default: 162)

SNMP Version 3

- ◆ **IP Address** – IPv4 or IPv6 address of a new management station to receive notification message (i.e., the targeted recipient).
- ◆ **Version** – Specifies whether to send notifications using SNMP v1, v2c, or v3 traps.
- ◆ **Notification Type**
 - **Traps** – Notifications are sent as trap messages.
 - **Inform** – Notifications are sent as inform messages. Note that this option is only available for version 2c and 3 hosts. (Default: traps are used)
 - **Timeout** – The number of seconds to wait for an acknowledgment before resending an inform message. (Range: 0-2147483647 centiseconds; Default: 1500 centiseconds)
 - **Retry times** – The maximum number of times to resend an inform message if the recipient does not acknowledge receipt. (Range: 0-255; Default: 3)
- ◆ **Local User Name** – The name of a local user which is used to identify the source of SNMPv3 trap messages sent from the local switch. (Range: 1-32 characters)

If an account for the specified user has not been created ([page 367](#)), one will be automatically generated.
- ◆ **Remote User Name** – The name of a remote user which is used to identify the source of SNMPv3 inform messages sent from the local switch. (Range: 1-32 characters)

If an account for the specified user has not been created ([page 369](#)), one will be automatically generated.
- ◆ **UDP Port** – Specifies the UDP port number used by the trap manager. (Default: 162)
- ◆ **Security Level** – When trap version 3 is selected, you must specify one of the following security levels. (Default: noAuthNoPriv)
 - **noAuthNoPriv** – There is no authentication or encryption used in SNMP communications.
 - **AuthNoPriv** – SNMP communications use authentication, but the data is not encrypted.
 - **AuthPriv** – SNMP communications use both authentication and encryption.

Web Interface

To configure trap managers:

1. Click Administration, SNMP.
2. Select Configure Trap from the Step list.
3. Select Add from the Action list.
4. Fill in the required parameters based on the selected SNMP version.
5. Click Apply

Figure 225: Configuring Trap Managers (SNMPv1)

The screenshot shows the 'Administration > SNMP' configuration page. At the top, there is a breadcrumb trail. Below it, the 'Step' is set to '6. Configure Trap' and the 'Action' is 'Add'. The form contains the following fields:

IP Address	192.168.0.3
Version	v1
Community String	private
UDP Port (1-65535)	162

At the bottom right, there are 'Apply' and 'Revert' buttons.

Figure 226: Configuring Trap Managers (SNMPv2c)

The screenshot shows the 'Administration > SNMP' configuration page for SNMPv2c. At the top, there is a breadcrumb trail. Below it, the 'Step' is set to '6. Configure Trap' and the 'Action' is 'Add'. The form contains the following fields:

IP Address	192.168.2.9	
Version	v2c	
Notification Type	Inform	
Timeout (0-2147483647)		centiseconds
Retry Times (0-255)		
Community String	venus	
UDP Port (1-65535)		

At the bottom right, there are 'Apply' and 'Revert' buttons.

Figure 227: Configuring Trap Managers (SNMPv3)

To show configured trap managers:

1. Click Administration, SNMP.
2. Select Configure Trap from the Step list.
3. Select Show from the Action list.

Figure 228: Showing Trap Managers

<input type="checkbox"/>	IP Address	Version	Community String/User Name	UDP Port	Security Level	Timeout	Retry Times
<input type="checkbox"/>	192.168.0.4	v3	steve	162	noAuthNoPriv		
<input type="checkbox"/>	192.168.0.5	v3	bobby	162	authNoPriv		
<input type="checkbox"/>	192.168.0.6	v3	betty	162	authNoPriv		

Creating SNMP Notification Logs

Use the Administration > SNMP (Configure Notify Filter - Add) page to create an SNMP notification log.

Command Usage

- ◆ Systems that support SNMP often need a mechanism for recording Notification information as a hedge against lost notifications, whether there are Traps or Informs that may be exceeding retransmission limits. The Notification Log MIB (NLM, RFC 3014) provides an infrastructure in which information from other MIBs may be logged.
- ◆ Given the service provided by the NLM, individual MIBs can now bear less responsibility to record transient information associated with an event against

the possibility that the Notification message is lost, and applications can poll the log to verify that they have not missed any important Notifications.

- ◆ If notification logging is not configured, when the switch reboots, some SNMP traps (such as warm start) cannot be logged.
- ◆ To avoid this problem, notification logging should be configured as described in this section, and these commands stored in the startup configuration file using the System > File (Copy – Running-Config) page as described on [page 79](#). Then when the switch reboots, SNMP traps (such as warm start) can now be logged.
- ◆ Based on the default settings used in RFC 3014, a notification log can contain up to 256 entries, and the entry aging time is 1440 minutes. Information recorded in a notification log, and the entry aging time can only be configured using SNMP from a network management station.
- ◆ When a trap host is created using the Administration > SNMP (Configure Trap – Add) page described on [page 372](#), a default notify filter will be created.

Parameters

These parameters are displayed:

- ◆ **IP Address** – The Internet address of a remote device. The specified target host must already have been configured using the Administration > SNMP (Configure Trap – Add) page.

The notification log is stored locally. It is not sent to a remote device. This remote host parameter is only required to complete mandatory fields in the SNMP Notification MIB.

- ◆ **Filter Profile Name** – Notification log profile name. (Range: 1-32 characters)

Web Interface

To create an SNMP notification log:

1. Click Administration, SNMP.
2. Select Configure Notify Filter from the Step list.
3. Select Add from the Action list.
4. Fill in the IP address of a configured trap manager and the filter profile name.
5. Click Apply

Figure 229: Creating SNMP Notification Logs

Administration > SNMP

Step: 7. Configure Notify Filter Action: Add

IP Address: 192.168.0.99

Filter Profile Name: R&D

Apply Revert

To show configured SNMP notification logs:

1. Click Administration, SNMP.
2. Select Configure Notify Filter from the Step list.
3. Select Show from the Action list.

Figure 230: Showing SNMP Notification Logs

Administration > SNMP

Step: 7. Configure Notify Filter Action: Show

SNMP Notify Filter List Total: 1

<input type="checkbox"/>	Filter profile name	IP Address
<input type="checkbox"/>	R&D	192.168.0.99

Delete Revert

Showing SNMP Statistics Use the Administration > SNMP (Show Statistics) page to show counters for SNMP input and output protocol data units.

Parameters

The following counters are displayed:

- ◆ **SNMP packets input** – The total number of messages delivered to the SNMP entity from the transport service.
- ◆ **Bad SNMP version errors** – The total number of SNMP messages which were delivered to the SNMP entity and were for an unsupported SNMP version.
- ◆ **Unknown community name** – The total number of SNMP messages delivered to the SNMP entity which used a SNMP community name not known to said entity.
- ◆ **Illegal operation for community name supplied** – The total number of SNMP messages delivered to the SNMP entity which represented an SNMP operation which was not allowed by the SNMP community named in the message.

- ◆ **Encoding errors** – The total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.
- ◆ **Number of requested variables** – The total number of MIB objects which have been retrieved successfully by the SNMP protocol entity as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
- ◆ **Number of altered variables** – The total number of MIB objects which have been altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.
- ◆ **Get-request PDUs** – The total number of SNMP Get-Request PDUs which have been accepted and processed, or generated, by the SNMP protocol entity.
- ◆ **Get-next PDUs** – The total number of SNMP Get-Next PDUs which have been accepted and processed, or generated, by the SNMP protocol entity.
- ◆ **Set-request PDUs** – The total number of SNMP Set-Request PDUs which have been accepted and processed, or generated, by the SNMP protocol entity.
- ◆ **SNMP packets output** – The total number of SNMP Messages which were passed from the SNMP protocol entity to the transport service.
- ◆ **Too big errors** – The total number of SNMP PDUs which were generated by the SNMP protocol entity and for which the value of the error-status field is “tooBig.”
- ◆ **No such name errors** – The total number of SNMP PDUs which were delivered to, or generated by, the SNMP protocol entity and for which the value of the error-status field is “noSuchName.”
- ◆ **Bad values errors** – The total number of SNMP PDUs which were delivered to, or generated by, the SNMP protocol entity and for which the value of the error-status field is “badValue.”
- ◆ **General errors** – The total number of SNMP PDUs which were delivered to, or generated by, the SNMP protocol entity and for which the value of the error-status field is “genErr.”
- ◆ **Response PDUs** – The total number of SNMP Get-Response PDUs which have been accepted and processed by, or generated by, the SNMP protocol entity.
- ◆ **Trap PDUs** – The total number of SNMP Trap PDUs which have been accepted and processed by, or generated by, the SNMP protocol entity.

To show SNMP statistics:

1. Click Administration, SNMP.
2. Select Show Statistics from the Step list.

Figure 231: Showing SNMP Statistics

Administration > SNMP			
Step: 8. Show Statistics			
SNMP Statistics			
SNMP packets Input	0	SNMP packets Output	0
Bad SNMP version errors	0	Too big errors	0
Unknown community name	0	No such name errors	0
Illegal operation for community name supplied	0	Bad values errors	0
Encoding errors	0	General errors	0
Number of requested variables	0	Response PDUs	0
Number of altered variables	0	Trap PDUs	0
Get-request PDUs	0		
Get-next PDUs	0		
Set-request PDUs	0		

Remote Monitoring

Remote Monitoring allows a remote device to collect information or respond to specified events on an independent basis. This switch is an RMON-capable device which can independently perform a wide range of tasks, significantly reducing network management traffic. It can continuously run diagnostics and log information on network performance. If an event is triggered, it can automatically notify the network administrator of a failure and provide historical information about the event. If it cannot connect to the management agent, it will continue to perform any specified tasks and pass data back to the management station the next time it is contacted.

The switch supports mini-RMON, which consists of the Statistics, History, Event and Alarm groups. When RMON is enabled, the system gradually builds up information about its physical interfaces, storing this information in the relevant RMON database group. A management agent then periodically communicates with the switch using the SNMP protocol. However, if the switch encounters a critical event, it can automatically send a trap message to the management agent which can then respond to the event if so configured.

Configuring RMON Alarms

Use the Administration > RMON (Configure Global - Add - Alarm) page to define specific criteria that will generate response events. Alarms can be set to test data over any specified time interval, and can monitor absolute or changing values (such as a statistical counter reaching a specific value, or a statistic changing by a certain amount over the set interval). Alarms can be set to respond to rising or falling thresholds. (However, note that after an alarm is triggered it will not be triggered again until the statistical value crosses the opposite bounding threshold and then back across the trigger threshold.

Command Usage

- ◆ If an alarm is already defined for an index, the entry must be deleted before any changes can be made.

Parameters

These parameters are displayed:

- ◆ **Index** – Index to this entry. (Range: 1-65535)
- ◆ **Variable** – The object identifier of the MIB variable to be sampled. Only variables of the type etherStatsEntry.n.n may be sampled.

Note that etherStatsEntry.n uniquely defines the MIB variable, and etherStatsEntry.n.n defines the MIB variable, plus the etherStatsIndex. For example, 1.3.6.1.2.1.16.1.1.6.1 denotes etherStatsBroadcastPkts, plus the etherStatsIndex of 1.
- ◆ **Interval** – The polling interval. (Range: 1-31622400 seconds)
- ◆ **Sample Type** – Tests for absolute or relative changes in the specified variable.
 - **Absolute** – The variable is compared directly to the thresholds at the end of the sampling period.
 - **Delta** – The last sample is subtracted from the current value and the difference is then compared to the thresholds.
- ◆ **Rising Threshold** – If the current value is greater than or equal to the rising threshold, and the last sample value was less than this threshold, then an alarm will be generated. After a rising event has been generated, another such event will not be generated until the sampled value has fallen below the rising threshold, reaches the falling threshold, and again moves back up to the rising threshold. (Range: 0-2147483647)
- ◆ **Rising Event Index** – The index of the event to use if an alarm is triggered by monitored variables reaching or crossing above the rising threshold. If there is no corresponding entry in the event control table, then no event will be generated. (Range: 0-65535)
- ◆ **Falling Threshold** – If the current value is less than or equal to the falling threshold, and the last sample value was greater than this threshold, then an alarm will be generated. After a falling event has been generated, another such event will not be generated until the sampled value has risen above the falling threshold, reaches the rising threshold, and again moves back down to the failing threshold. (Range: 0-2147483647)
- ◆ **Falling Event Index** – The index of the event to use if an alarm is triggered by monitored variables reaching or crossing below the falling threshold. If there is no corresponding entry in the event control table, then no event will be generated. (Range: 0-65535)
- ◆ **Owner** – Name of the person who created this entry. (Range: 1-32 characters)

Web Interface

To configure an RMON alarm:

1. Click Administration, RMON.
2. Select Configure Global from the Step list.
3. Select Add from the Action list.
4. Click Alarm.
5. Enter an index number, the MIB object to be polled (etherStatsEntry.n.n), the polling interval, the sample type, the thresholds, and the event to trigger.
6. Click Apply

Figure 232: Configuring an RMON Alarm

The screenshot shows the 'Administration > RMON' configuration page. At the top, there are two dropdown menus: 'Step: 1. Configure Global' and 'Action: Add'. Below these are two radio buttons: 'Alarm' (selected) and 'Event'. The main configuration area contains several input fields and a dropdown menu:

Index (1-65535)	1
Variable	6.1
Interval (1-31622400)	15 sec
Sample Type	Delta
Rising Threshold (0-2147483647)	100
Rising Event Index (0-65535)	30
Falling Threshold (0-2147483647)	1
Falling Event Index (0-65535)	2
Owner	bill

At the bottom right of the form are two buttons: 'Apply' and 'Revert'.

To show configured RMON alarms:

1. Click Administration, RMON.
2. Select Configure Global from the Step list.
3. Select Show from the Action list.
4. Click Alarm.

Figure 233: Showing Configured RMON Alarms

The screenshot shows the 'Administration > RMON' configuration page. At the top, there is a breadcrumb trail and a 'Step: 1. Configure Global' dropdown with an 'Action: Show' dropdown. Below this, there are radio buttons for 'Alarm' (selected) and 'Event'. The main content is a table titled 'RMON Alarm List' with a total of 28 entries. The table has columns for Index, Status, Variable, Interval, Type, Last Value, Rising Threshold, Rising Event Index, Falling Threshold, Falling Event Index, and Owner. Five entries are visible, all with a status of 'Valid' and a type of 'Delta'.

<input type="checkbox"/>	Index	Status	Variable	Interval	Type	Last Value	Rising Threshold	Rising Event Index	Falling Threshold	Falling Event Index	Owner
<input type="checkbox"/>	1	Valid	1.3.6.1.2.1.16.1.1.1.6.1	30	Delta	0	892800	0	446400	0	
<input type="checkbox"/>	2	Valid	1.3.6.1.2.1.16.1.1.1.6.2	30	Delta	0	892800	0	446400	0	
<input type="checkbox"/>	3	Valid	1.3.6.1.2.1.16.1.1.1.6.3	30	Delta	0	892800	0	446400	0	
<input type="checkbox"/>	4	Valid	1.3.6.1.2.1.16.1.1.1.6.4	30	Delta	0	892800	0	446400	0	
<input type="checkbox"/>	5	Valid	1.3.6.1.2.1.16.1.1.1.6.5	30	Delta	0	892800	0	446400	0	

Configuring RMON Events

Use the Administration > RMON (Configure Global - Add - Event) page to set the action to take when an alarm is triggered. The response can include logging the alarm or sending a message to a trap manager. Alarms and corresponding events provide a way of immediately responding to critical network problems.

Command Usage

- ◆ If an alarm is already defined for an index, the entry must be deleted before any changes can be made.
- ◆ One default event is configured as follows:
 - event Index = 1
 - Description: RMON_TRAP_LOG
 - Event type: log & trap
 - Event community name is public
 - Owner is RMON_SNMP

Parameters

These parameters are displayed:

- ◆ **Index** – Index to this entry. (Range: 1-65535)
- ◆ **Type** – Specifies the type of event to initiate:
 - **None** – No event is generated.
 - **Log** – Generates an RMON log entry when the event is triggered. Log messages are processed based on the current configuration settings for event logging (see “System Log Configuration” on page 327).
 - **Trap** – Sends a trap message to all configured trap managers (see “Specifying Trap Managers” on page 372).
 - **Log and Trap** – Logs the event and sends a trap message.

- ◆ **Community** – A password-like community string sent with the trap operation to SNMP v1 and v2c hosts.

Although the community string can be set on this configuration page, it is recommended that it be defined on the SNMP trap configuration page (see [“Setting Community Access Strings” on page 366](#)) prior to configuring it here. (Range: 1-32 characters)

- ◆ **Description** – A comment that describes this event. (Range: 1-127 characters)

- ◆ **Owner** – Name of the person who created this entry. (Range: 1-32 characters)

Web Interface

To configure an RMON event:

1. Click Administration, RMON.
2. Select Configure Global from the Step list.
3. Select Add from the Action list.
4. Click Event.
5. Enter an index number, the type of event to initiate, the community string to send with trap messages, the name of the person who created this event, and a brief description of the event.
6. Click Apply

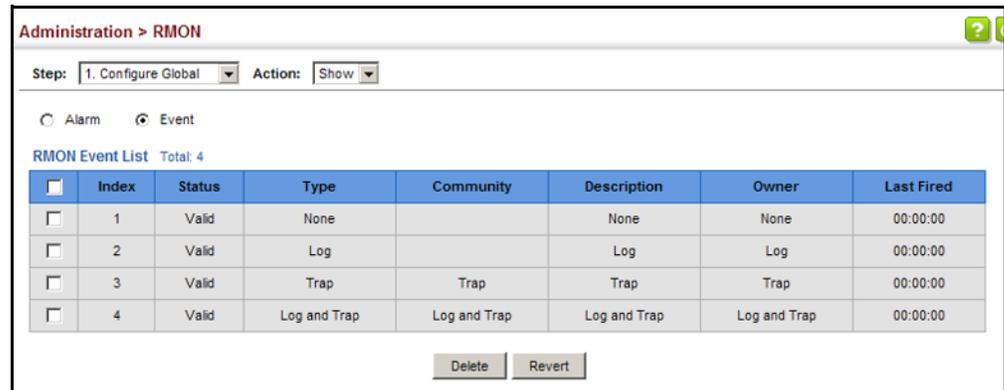
Figure 234: Configuring an RMON Event

The screenshot shows a web interface for configuring an RMON event. The breadcrumb path is "Administration > RMON". At the top, there are two dropdown menus: "Step:" set to "1. Configure Global" and "Action:" set to "Add". Below these are two radio buttons: "Alarm" (unselected) and "Event" (selected). The main configuration area contains several fields: "Index (1-65535)" with the value "2", "Type" with a dropdown menu set to "Log and Trap", "Community" with the text "private", "Description" with the text "for software group", and "Owner" with the text "david". At the bottom right, there are two buttons: "Apply" and "Revert".

To show configured RMON events:

1. Click Administration, RMON.
2. Select Configure Global from the Step list.
3. Select Show from the Action list.
4. Click Event.

Figure 235: Showing Configured RMON Events



Configuring RMON History Samples

Use the Administration > RMON (Configure Interface - Add - History) page to collect statistics on a physical interface to monitor network utilization, packet types, and errors. A historical record of activity can be used to track down intermittent problems. The record can be used to establish normal baseline activity, which may reveal problems associated with high traffic levels, broadcast storms, or other unusual events. It can also be used to predict network growth and plan for expansion before your network becomes too overloaded.

Command Usage

- ◆ Each index number equates to a port on the switch.
- ◆ If history collection is already enabled on an interface, the entry must be deleted before any changes can be made.
- ◆ The information collected for each sample includes:
input octets, packets, broadcast packets, multicast packets, undersize packets, oversize packets, fragments, jabbers, CRC alignment errors, collisions, drop events, and network utilization.
For a description of the statistics displayed on the Show Details page, refer to [“Showing Port or Trunk Statistics” on page 114](#).
- ◆ The switch reserves two index entries for each port. If a default index entry is re-assigned to another port using the Add page, this index will not appear in the Show nor Show Details page for the port to which is normally assigned. For

example, if control entry 15 is assigned to port 5, this index entry will be removed from the Show and Show Details page for port 8.

Parameters

These parameters are displayed:

- ◆ **Port** – The port number on the switch.
- ◆ **Index** – Index to this entry. (Range: 1-65535)
- ◆ **Interval** - The polling interval. (Range: 1-3600 seconds; Default: 1800 seconds)
- ◆ **Buckets** – The number of buckets requested for this entry. (Range: 1-65536; Default: 8)
The number of buckets granted are displayed on the Show page.
- ◆ **Owner** - Name of the person who created this entry. (Range: 1-32 characters)

Web Interface

To periodically sample statistics on a port:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Add from the Action list.
4. Click History.
5. Select a port from the list as the data source.
6. Enter an index number, the sampling interval, the number of buckets to use, and the name of the owner for this entry.
7. Click Apply

Figure 236: Configuring an RMON History Sample

Administration > RMON

Step: 2. Configure Interface Action: Add

History Statistics

Port 2

Index (1-65535) 100

Interval (1-3600) 60 sec

Buckets (1-65535) 10

Owner david

Apply Revert

To show configured RMON history samples:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Show from the Action list.
4. Select a port from the list.
5. Click History.

Figure 237: Showing Configured RMON History Samples

Administration > RMON

Step: 2. Configure Interface Action: Show

History Statistics

Port 1

RMON History Port List Total: 2

<input type="checkbox"/>	Index	Status	Interval	Requested Buckets	Granted Buckets	Owner
<input type="checkbox"/>	1	Valid	1800	50	8	
<input type="checkbox"/>	2	Valid	30	50	8	

Delete Revert

To show collected RMON history samples:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Show Details from the Action list.
4. Select a port from the list.

5. Click History.

Figure 238: Showing Collected RMON History Samples

History Index	Sample Index	Interval Start	Octets	Packets	Broadcast Packets	Multicast Packets	Undersize Packets	Oversize Packets	Fragments	Jabbers	CRC Align Errors	Collisions	Drop Events	Network Utilization
1	1	00:00:00	1735989	4434	20	67	0	0	0	0	0	0	0	0
2	94	00:46:30	12870	43	0	1	0	0	0	0	0	0	0	0
2	95	00:47:00	19724	61	0	1	0	0	0	0	0	0	0	0
2	96	00:47:30	26146	71	0	1	0	0	0	0	0	0	0	0
2	97	00:48:00	22012	60	0	1	0	0	0	0	0	0	0	0

Configuring RMON Statistical Samples

Use the Administration > RMON (Configure Interface - Add - Statistics) page to collect statistics on a port, which can subsequently be used to monitor the network for common errors and overall traffic rates.

Command Usage

- ◆ If statistics collection is already enabled on an interface, the entry must be deleted before any changes can be made.
- ◆ The information collected for each entry includes:
 - input octets, packets, broadcast packets, multicast packets, undersize packets, oversize packets, CRC alignment errors, jabbers, fragments, collisions, drop events, and frames of various sizes.

Parameters

These parameters are displayed:

- ◆ **Port** – The port number on the switch.
- ◆ **Index** – Index to this entry. (Range: 1-65535)
- ◆ **Owner** – Name of the person who created this entry. (Range: 1-32 characters)

Web Interface

To enable regular sampling of statistics on a port:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Add from the Action list.
4. Click Statistics.

5. Select a port from the list as the data source.
6. Enter an index number, and the name of the owner for this entry
7. Click Apply

Figure 239: Configuring an RMON Statistical Sample

The screenshot shows the 'Administration > RMON' configuration page. At the top, the breadcrumb is 'Administration > RMON'. Below it, the 'Step' is set to '2. Configure Interface' and the 'Action' is 'Add'. There are two radio buttons: 'History' (unselected) and 'Statistics' (selected). The 'Port' is set to '2'. The 'Index (1-65535)' field contains '100' and the 'Owner' field contains 'mary'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show configured RMON statistical samples:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Show from the Action list.
4. Select a port from the list.
5. Click Statistics.

Figure 240: Showing Configured RMON Statistical Samples

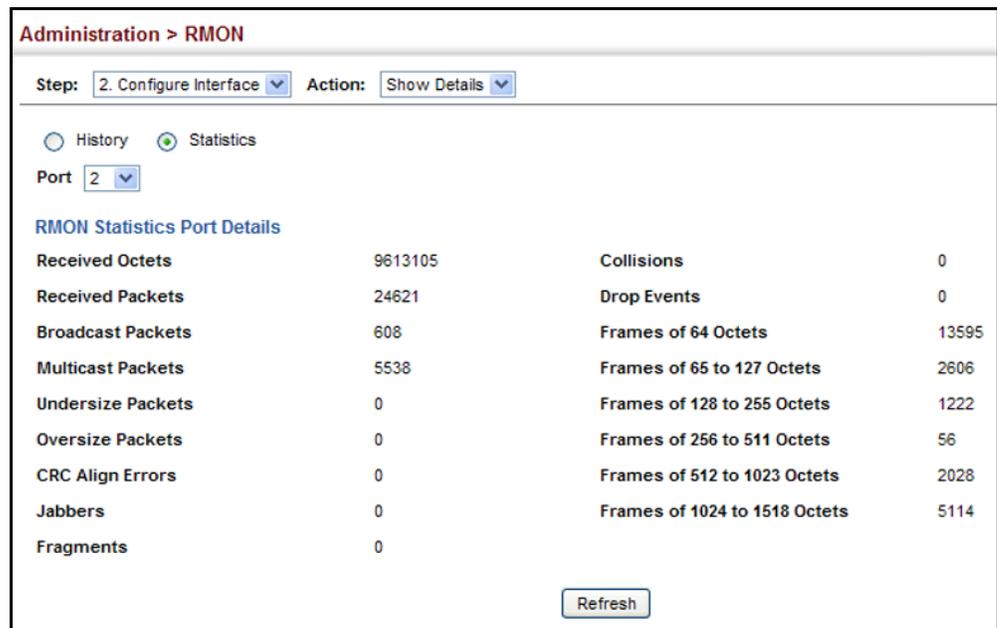
The screenshot shows the 'Administration > RMON' configuration page. At the top, the breadcrumb is 'Administration > RMON'. Below it, the 'Step' is '2. Configure Interface' and the 'Action' is 'Show'. There are two radio buttons: 'History' (unselected) and 'Statistics' (selected). The 'Port' is set to '2'. Below this, there is a table titled 'RMON Statistics Port List' with a total of 2 entries. The table has columns for 'Index', 'Status', and 'Owner'. At the bottom right, there are 'Delete' and 'Revert' buttons.

	Index	Status	Owner
<input type="checkbox"/>	1	Valid	abc
<input type="checkbox"/>	2	Valid	test

To show collected RMON statistical samples:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Show Details from the Action list.
4. Select a port from the list.
5. Click Statistics.

Figure 241: Showing Collected RMON Statistical Samples



Connectivity Fault Management

Connectivity Fault Management (CFM) is an OAM protocol that includes proactive connectivity monitoring using continuity check messages, fault verification through loop back messages, and fault isolation by examining end-to-end connections between provider edge devices or between customer edge devices.

CFM is implemented as a service level protocol based on service instances which encompass only that portion of the metropolitan area network supporting a specific customer. CFM can also provide controlled management access to a hierarchy of maintenance domains (such as the customer, service provider, and equipment operator).

This switch supports functions for defining the CFM structure, including domains, maintenance associations, and maintenance access points. It also supports fault detection through continuity check messages for all known maintenance points,

and cross-check messages which are used to verify a static list of remote maintenance points located on other devices (in the same maintenance association) against those found through continuity check messages. Fault verification is supported using loop back messages, and fault isolation with link trace messages. Fault notification is also provided by SNMP alarms which are automatically generated by maintenance points when connectivity faults or configuration errors are detected in the local maintenance domain.

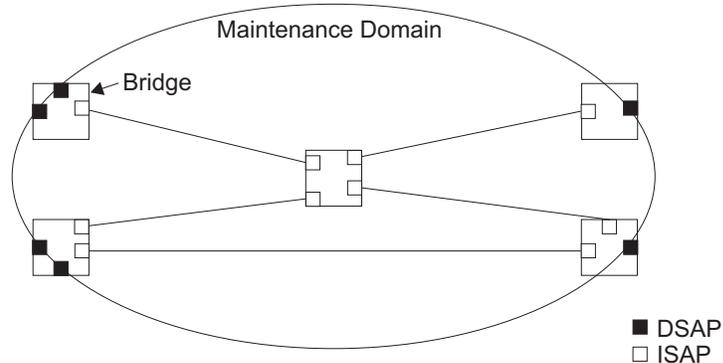
Key Components of CFM

CFM provides restricted management access to each Service Instance using a structured conceptual network based on these components:

- ◆ A Maintenance Domain defines a part of the network controlled by a single operator, and supports management access to the domain through Domain Service Access Points (DSAPs) configured on the domain boundary, as well as connectivity testing between these DSAPs.
- ◆ A Maintenance Association (MA) contains the DSAPs for an individual Service Instance. DSAPs are the primary maintenance points used to monitor connectivity across a maintenance domain, and are the entry points to the paths which interconnect the access points allocated to a service instance.
- ◆ A Maintenance Level allows maintenance domains to be nested in a hierarchical fashion, providing access to the specific network portions required by each operator. Domains at lower levels may be either hidden or exposed to operators managing domains at a higher level, allowing either course or fine fault resolution.
- ◆ Maintenance End Points (MEPs) which provide full CFM access to a Service Instance (i.e., a specific MA), and Maintenance Intermediate Points (MIPs) which are passive entities that merely validate received CFM messages, or respond to link trace and loop back requests. MIPs are the interconnection points that make up all possible paths between the DSAPs within an MA, and may also include interconnection points in lower-level domains if exposed by CFM settings.

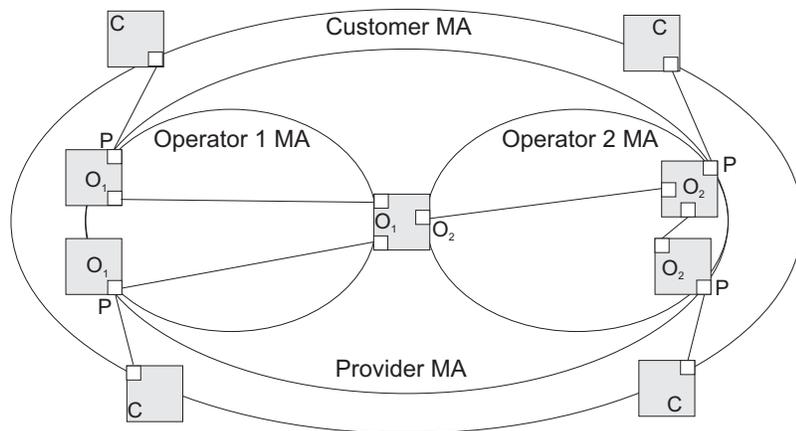
The following figure shows a single Maintenance Domain, with DSAPs located on the domain boundary, and Internal Service Access Points (ISAPs) inside the domain through which frames may pass between the DSAPs.

Figure 242: Single CFM Maintenance Domain



The figure below shows four maintenance associations contained within a hierarchical structure of maintenance domains. At the innermost level, there are two operator domains which include access points marked "O₁" and "O₂" respectively. The users of these domains can see their respective MEPs as well as all the MIPs within their domains. There is a service provider domain at the second level in the hierarchy. From the service provider's view, the access points marked "P" are visible, and all access points within the operator domains have also been made visible as MIPs according to common practice. And finally, there is a customer domain at the top of the hierarchy. Users at this level can only see the access points marked "C" on the outer domain boundary. Again, normal practice is to hide the internal structure of the network from outsiders to reduce security risks.

Figure 243: Multiple CFM Maintenance Domains



Note that the Service Instances within each domain shown above are based on a unique maintenance association for the specific users, distinguished by the domain name, maintenance level, maintenance association's name, and assigned VLAN.

Basic CFM Operations

CFM uses standard Ethernet frames for sending protocol messages. Both the source and destination address for these messages are based on unicast or multicast MAC addresses, and therefore confined to a single Layer 2 CFM service VLAN. For this reason, the transmission, forwarding, and processing of CFM frames is performed by bridges, not routers. Bridges that do not recognize CFM messages forward them as normal data. There are three basic types of CFM messages, including continuity check, link trace, and loop back.

Continuity check messages (CCMs) are multicast within a single Service Instance (i.e., a specific MA), allowing MEPs to discover other MEPs within the same MA, and MIPs to discover MEPs. Connectivity faults are indicated when a known MEP stops sending CCMs, or a remote MEP configured in a static list does not come up. Configuration errors, such as a cross-connect between different MAs, are indicated when a CCM is received with an incorrect MA identifier or maintenance level.

Loop back messages are used for fault verification. These messages can be sent using the MAC address of any destination MEP within the same MA. If the target MEP's identifier has been discovered through CCM messages, then a loop back message can also be sent using the MEPs identifier. A reply indicates that the destination is reachable.

Link trace messages are used for fault verification. These messages are multicast frames sent out to track the hop-by-hop path to a target MEP within the same MA. Responses provide information on the ingress, egress, and relay action taken at each hop along the path, providing vital information about connectivity problems. Responses allow the sender to discover all of the maintenance points that would be traversed by a data frame sent to the target MAC address.

SNMP traps can also be configured to provide an automated method of fault notification. If the fault notification generator detects one or more defects within the configured time period, and fault alarms are enabled, a corresponding trap will be sent. No further fault alarms are sent until the fault notification generator has been reset by the passage of a configured time period without detecting any further faults. Upon receiving a fault alarm, you should inspect the related SNMP objects for the reporting MEP, diagnose the fault, correct it, and re-examine the MEP's SNMP objects to see whether the fault notification generator has been reset.

Configuration Guidelines

1. Configure the maintenance domains with the MD List (see "[Configuring CFM Maintenance Domains](#)").
2. Configure the maintenance associations with MA List (see "[Configuring CFM Maintenance Associations](#)").
3. Configure the local maintenance end points (MEPs) which will serve as the domain service access points for the specified maintenance association using the MEP List (see "[Configuring CFM Maintenance Associations](#)").

4. Enter a static list of MEPs assigned to other devices within the same maintenance association using the Remote MEP List (see ["Configuring Remote Maintenance End Points"](#)). This allows CFM to automatically verify the functionality of these remote end points by cross-checking the static list configured on this device against information learned through continuity check messages.
5. Enable CFM globally on the switch using the Configure Global screen (see ["Configuring Global Settings for CFM"](#)).
6. Enable CFM on the local MEPs using the Configure Interface screen (see ["Configuring Interfaces for CFM"](#)).
7. Enable continuity check and cross-check operations, and configure AIS parameters using the Configure MA – Configure Details screen (see ["Configuring CFM Maintenance Associations"](#)).

Other configuration changes may be required for your particular environment, such as adjusting the interval at which continuity check messages are sent (see ["Configuring CFM Maintenance Associations"](#)), or setting the start-up delay for the cross-check operation (see ["Configuring Global Settings for CFM"](#)). You can also enable SNMP traps for events discovered by continuity check messages or cross-check messages (see ["Configuring Global Settings for CFM"](#)).

Configuring Global Settings for CFM

Use the Administration > CFM (Configure Global) page to configure global settings for CFM, such as enabling the CFM process on the switch, setting the start-up delay for cross-check operations, configuring parameters for the link trace cache, and enabling traps for events discovered by continuity check messages or cross-check messages.

Parameters

These parameters are displayed:

Global Configuration

- ◆ **CFM Status** – Enables CFM processing globally on the switch. (Default: Enabled)

To avoid generating an excessive number of traps, the complete CFM maintenance structure and process parameters should be configured prior to enabling CFM processing globally on the switch. Specifically, the maintenance domains, maintenance associations, and maintenance end-points (MEPs) should be configured on each participating bridge using the Configure MD page (see ["Configuring CFM Maintenance Domains"](#)), Configure MA page (see ["Configuring CFM Maintenance Associations"](#)), and the Configure MEP page (see ["Configuring Maintenance End Points"](#)).

When CFM is enabled, hardware resources are allocated for CFM processing.

- ◆ **MEP Cross Check Start Delay** – Sets the maximum delay that a device waits for remote MEPs to come up before starting the cross-check operation. (Range: 1-65535 seconds; Default: 10 seconds)

This parameter sets the time to wait for a remote MEP to come up, and the switch starts cross-checking the list of statically configured remote MEPs in the local maintenance domain (Configure Remote MEP page, see "[Configuring Remote Maintenance End Points](#)") against the MEPs learned through continuity check messages (CCMs).

The cross-check start delay should be configured to a value greater than or equal to the continuity check message interval to avoid generating unnecessary traps (see "[Configuring CFM Maintenance Associations](#)").

Link Trace Cache Settings

- ◆ **Link Trace Cache** – Enables caching of CFM data learned through link trace messages. (Default: Enabled)

A linktrace message is a multicast CFM frame initiated by a MEP, and forwarded from MIP to MIP, with each MIP generating a linktrace reply, up to the point at which the linktrace message reaches its destination or can no longer be forwarded.

Use this command attribute to enable the link trace cache to store the results of link trace operations initiated on this device. Use the CFM Transmit Link Trace page (see "[Transmitting Link Trace Messages](#)") to transmit a linktrace message.

Linktrace responses are returned from each MIP along the path and from the target MEP. Information stored in the cache includes the maintenance domain name, MA name, MEPID, sequence number, and TTL value (see "[Displaying Fault Notification Settings](#)").

- ◆ **Link Trace Cache Hold Time** – The hold time for CFM link trace cache entries. (Range: 1-65535 minutes; Default: 100 minutes)

Before setting the aging time for cache entries, the cache must first be enabled in the Linktrace Cache attribute field.

- ◆ **Link Trace Cache Size** – The maximum size for the link trace cache. (Range: 1-4095 entries; Default: 100 entries)

If the cache reaches the maximum number of specified entries, or the size is set to a value less than the current number of stored entries, no new entries are added. To add additional entries, the cache size must first be increased, or purged (see "[Displaying Fault Notification Settings](#)").

Continuity Check Errors

- ◆ **Connectivity Check Config** – Sends a trap if this device receives a continuity check message (CCM) with the same maintenance end point identifier (MPID) as its own but with a different source MAC address, indicating that a CFM configuration error exists.

- ◆ **Connectivity Check Loop** – Sends a trap if this device receives a CCM with the same source MAC address and MPID as its own, indicating that a forwarding loop exists.
- ◆ **Connectivity Check MEP Down** – Sends a trap if this device loses connectivity with a remote maintenance end point (MEP), or connectivity has been restored to a remote MEP which has recovered from an error condition.
- ◆ **Connectivity Check MEP Up** – Sends a trap if a remote MEP is discovered and added to the local database, the port state of a previously discovered remote MEP changes, or a CCM is received from a remote MEP which as an expired entry in the archived database.

MEP Up traps are suppressed when cross-checking of MEPs is enabled¹¹ because cross-check traps include more detailed status information.

Cross-check Errors

- ◆ **Cross Check MA Up** – Sends a trap when all remote MEPs in an MA come up.
An MA Up trap is sent if cross-checking is enabled¹¹, and a CCM is received from all remote MEPs configured in the static list for this maintenance association¹².
- ◆ **Cross Check MEP Missing** – Sends a trap if the cross-check timer expires and no CCMs have been received from a remote MEP configured in the static list.
A MEP Missing trap is sent if cross-checking is enabled¹¹, and no CCM is received for a remote MEP configured in the static list¹².
- ◆ **Cross Check MEP Unknown** – Sends a trap if an unconfigured MEP comes up.
A MEP Unknown trap is sent if cross-checking is enabled¹¹, and a CCM is received from a remote MEP that is not configured in the static list¹².

Web Interface

To configure global settings for CFM:

1. Click Administration, CFM.
2. Select Configure Global from the Step list.
3. Before enabling CFM processing on the switch, first configure the required CFM domains, maintenance associations, and static MEPs. Then set the delay time to wait for a remote MEP comes up before the switch starts cross-checking the end points learned through CCMs against those stored in the static list.
4. Adjust the parameters for the link trace cache as required.

11. Cross-checking must be enabled for this type of trap to be reported (see "[Configuring CFM Maintenance Associations](#)").

12. See "[Configuring Maintenance End Points](#)".

5. Enable the required traps for continuity check and cross-check errors. Remember that the “Connectivity Check” and “Cross Check” fields on the MA Configuration page must be enabled before related errors can be generated.
6. Click Apply.

Figure 244: Configuring Global Settings for CFM

The screenshot shows the 'Administration > CFM' configuration page. At the top, there is a 'Step:' dropdown menu set to '1. Configure Global'. Below this, the page is divided into two main sections: 'Global Configuration' and 'SNMP Trap Configuration'.
Global Configuration:
 - CFM Status: Enabled
 - MEP Cross Check Start Delay (1-65535): sec
 - Link Trace Cache: Enabled
 - Link Trace Cache Hold Time (1-65535): min
 - Link Trace Cache Size (1-4095): entries
SNMP Trap Configuration:
 - Connectivity Check Config: Enabled
 - Connectivity Check Loop: Enabled
 - Connectivity Check MEP Down: Enabled
 - Connectivity Check MEP Up: Enabled
 - Cross Check MA Up: Enabled
 - Cross Check MEP Missing: Enabled
 - Cross Check MEP Unknown: Enabled
 At the bottom right of the configuration area, there are two buttons: 'Apply' and 'Revert'.

Configuring Interfaces for CFM CFM processes are enabled by default for all physical interfaces, both ports and trunks. You can use the Administration > CFM (Configure Interface) page to change these settings.

Command Usage

- ◆ An interface must be enabled before a MEP can be created (see "[Configuring Maintenance End Points](#)").
- ◆ If a MEP has been configured on an interface, it must first be deleted before CFM can be disabled on that interface.
- ◆ When CFM is disabled, hardware resources previously used for CFM processing on that interface are released, and all CFM frames entering that interface are forwarded as normal data traffic.

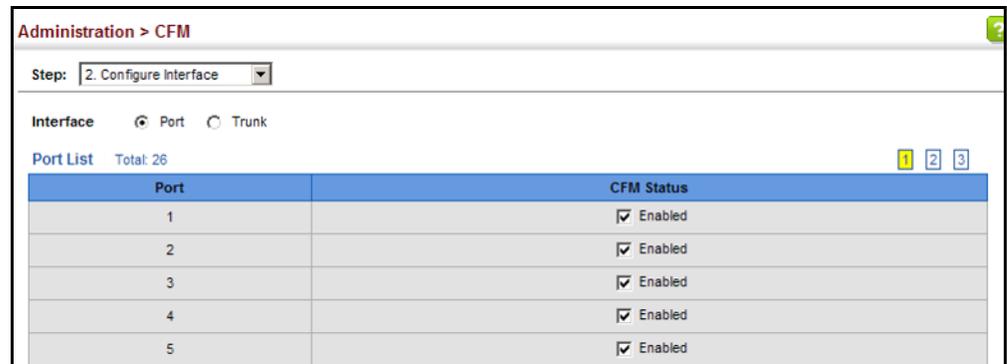
Web Interface

To enable CFM on an interface:

1. Click Administration, CFM.

2. Select Configure Interface from the Step list.
3. Select Port or Trunk.
4. Enable CFM on the required interface.
5. Click Apply.

Figure 245: Configuring Interfaces for CFM



Port	CFM Status
1	✓ Enabled
2	✓ Enabled
3	✓ Enabled
4	✓ Enabled
5	✓ Enabled

Configuring CFM Maintenance Domains

Use the Administration > CFM (Configure MD) pages to create and configure a Maintenance Domain (MD) which defines a portion of the network for which connectivity faults can be managed. Domain access points are set up on the boundary of a domain to provide end-to-end connectivity fault detection, analysis, and recovery. Domains can be configured in a hierarchy to provide management access to the same basic network resources for different user levels.

Command Usage

Configuring General Settings

- ◆ Where domains are nested, an upper-level hierarchical domain must have a higher maintenance level than the ones it encompasses. The higher to lower level domain types commonly include entities such as customer, service provider, and operator.
- ◆ More than one domain can be configured at the same maintenance level, but a single domain can only be configured with one maintenance level.
- ◆ If MEPs (see "[Configuring Maintenance End Points](#)") or MAs (see "[Configuring CFM Maintenance Associations](#)") are configured for a domain, they must first be removed before you can remove the domain.

Maintenance domains are designed to provide a transparent method of verifying and resolving connectivity problems for end-to-end connections. By default, these connections run between the domain service access points (DSAPs) within each MA defined for a domain, and are manually configured (see "[Configuring Maintenance End Points](#)").

In contrast, MIPs are interconnection points that make up all possible paths between the DSAPs within an MA. MIPs are automatically generated by the CFM protocol when the MIP Creation Type is set to "Default" or "Explicit," and the MIP creation state machine is invoked (as defined in IEEE 802.1ag). The default option allows MIPs to be created for all interconnection points within an MA, regardless of the domain's level in the maintenance hierarchy (e.g., customer, provider, or operator). While the explicit option only generates MIPs within an MA if its associated domain is not at the bottom of the maintenance hierarchy. This option is used to hide the structure of network at the lowest domain level.

The diagnostic functions provided by CFM can be used to detect connectivity failures between any pair of MEPs in an MA. Using MIPs allows these failures to be isolated to smaller segments of the network.

Allowing the CFM to generate MIPs exposes more of the network structure to users at higher domain levels, but can speed up the process of fault detection and recovery. This trade-off should be carefully considered when designing a CFM maintenance structure.

Also note that while MEPs are active agents which can initiate consistency check messages (CCMs), transmit loop back or link trace messages, and maintain the local CCM database, MIPs, on the other hand, are passive agents which can only validate received CFM messages, and respond to loop back and link trace messages.

The MIP creation method defined for an MA (see "[Configuring CFM Maintenance Associations](#)") takes precedence over the method defined on the CFM Domain List.

Configuring Fault Notification

- ◆ A fault alarm can generate an SNMP notification. It is issued when the MEP fault notification generator state machine detects that the configured time period (MEP Fault Notify Alarm Time) has passed with one or more defects indicated, and fault alarms are enabled at or above the specified priority level (MEP Fault Notify Lowest Priority). The state machine transmits no further fault alarms until it is reset by the passage of a configured time period (MEP Fault Notify Reset Time) without a defect indication. The normal procedure upon receiving a fault alarm is to inspect the reporting MEP's managed objects using an appropriate SNMP software tool, diagnose the fault, correct it, re-examine the MEP's managed objects to see whether the MEP fault notification generator state machine has been reset, and repeat those steps until the fault is resolved.
- ◆ Only the highest priority defect currently detected is reported in the fault alarm.

Priority levels include the following options:

Table 34: Remote MEP Priority Levels

Priority Level	Level Name	Description
1	allDef	All defects.
2	macRemErrXcon	DefMACstatus, DefRemoteCCM, DefErrorCCM, or DefXconCCM.
3	remErrXcon	DefErrorCCM, DefXconCCM or DefRemoteCCM.
4	errXcon	DefErrorCCM or DefXconCCM.
5	xcon	DefXconCCM
6	noXcon	No defects DefXconCCM or lower are to be reported.

Table 35: MEP Defect Descriptions

Defect	Description
DefMACstatus	Either some remote MEP is reporting its Interface Status TLV as not isUp, or all remote MEPs are reporting a Port Status TLV that contains some value other than psUp.
DefRemoteCCM	The MEP is not receiving valid CCMs from at least one of the remote MEPs.
DefErrorCCM	The MEP has received at least one invalid CCM whose CCM Interval has not yet timed out.
DefXconCCM	The MEP has received at least one CCM from either another MAID or a lower MD Level whose CCM Interval has not yet timed out.

Parameters

These parameters are displayed:

Creating a Maintenance Domain

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MD Name** – Maintenance domain name. (Range: 1-43 alphanumeric characters)
- ◆ **MD Level** – Authorized maintenance level for this domain. (Range: 0-7)
- ◆ **MIP Creation Type** – Specifies the CFM protocol's creation method for maintenance intermediate points (MIPs) in this domain:
 - **Default** – MIPs can be created for any maintenance association (MA) configured in this domain on any bridge port through which the MA's VID can pass.
 - **Explicit** – MIPs can be created for any MA configured in this domain only on bridge ports through which the MA's VID can pass, and only if a maintenance end point (MEP) is created at some lower MA Level.
 - **None** – No MIP can be created for any MA configured in this domain.

Configuring Detailed Settings for a Maintenance Domain

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MEP Archive Hold Time** – The time that data from a missing MEP is retained in the continuity check message (CCM) database before being purged. (Range: 1-65535 minutes; Default: 100 minutes)
A change to the hold time only applies to entries stored in the database after this attribute is changed.
- ◆ **MEP Fault Notify Lowest Priority** – The lowest priority defect that is allowed to generate a fault alarm. (Range: 1-6, Default: 2)
- ◆ **MEP Fault Notify Alarm Time** – The time that one or more defects must be present before a fault alarm is issued. (Range: 3-10 seconds; Default: 3 seconds)
- ◆ **MEP Fault Notify Reset Time** – The time after a fault alarm has been issued, and no defect exists, before another fault alarm can be issued. (Range: 3-10 seconds; Default: 10 seconds)

Web Interface

To create a maintenance domain:

1. Click Administration, CFM.
2. Select Configure MD from the Step list.
3. Select Add from the Action list.
4. Specify the maintenance domains and authorized maintenance levels (thereby setting the hierarchical relationship with other domains).
5. Specify the manner in which MIPs can be created within each domain.
6. Click Apply.

Figure 246: Configuring Maintenance Domains

Administration > CFM

Step: 1. Configure MD Action: Add

MD Index (1-65535) 1

MD Name voip

MD Level (0-7) 3

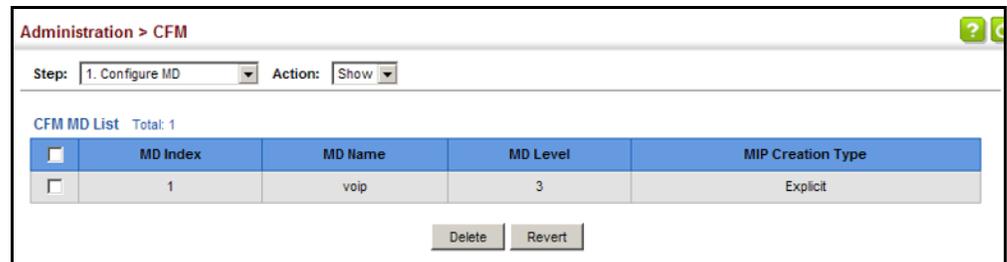
MIP Creation Type Explicit

Apply Revert

To show the configured maintenance domains:

1. Click Administration, CFM.
2. Select Configure MD from the Step list.
3. Select Show from the Action list.

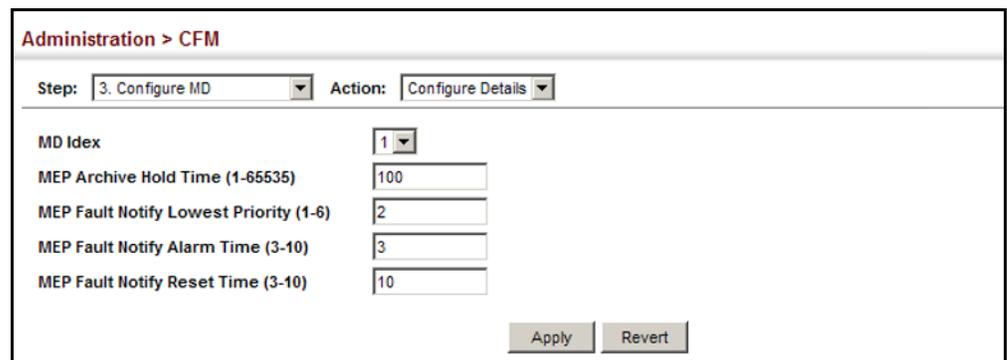
Figure 247: Showing Maintenance Domains



To configure detailed settings for maintenance domains:

1. Click Administration, CFM.
2. Select Configure MD from the Step list.
3. Select Configure Details from the Action list.
4. Select an entry from the MD Index.
5. Specify the MEP archive hold and MEP fault notification parameters.
6. Click Apply

Figure 248: Configuring Detailed Settings for Maintenance Domains



Configuring CFM Maintenance Associations Use the Administration > CFM (Configure MA) pages to create and configure the Maintenance Associations (MA) which define a unique CFM service instance. Each MA can be identified by its parent MD, the MD's maintenance level, the VLAN assigned to the MA, and the set of maintenance end points (MEPs) assigned to it.

Command Usage

Creating a Maintenance Association

- ◆ Use the Configure MA – Add screen to create an MA within the selected MD, map it to a customer service instance (S-VLAN), and set the manner in which MIPs are created for this service instance. Then use the MEP List to assign domain service access points (DSAPs) to this service instance (see [“Configuring Maintenance End Points” on page 407](#)).
- ◆ An MA must be defined before any associated DSAPs or remote MEPs can be assigned (see [“Configuring Remote Maintenance End Points” on page 409](#)).
- ◆ Multiple domains at the same maintenance level cannot have an MA on the same VLAN (see [“Configuring CFM Maintenance Domains” on page 398](#)).
- ◆ Before removing an MA, first remove the MEPs assigned to it (see [“Configuring Maintenance End Points” on page 407](#)).
- ◆ For a detailed description of the MIP types, refer to the Command Usage section under [“Configuring CFM Maintenance Domains” on page 398](#).

Configuring Detailed Settings for a Maintenance Association

- ◆ CCMs are multicast periodically by a MEP in order to discover other MEPs in the same MA, and to assure connectivity to all other MEPs/MIPs in the MA.
- ◆ Each CCM received is checked to verify that the MEP identifier field sent in the message does not match its own MEP ID, which would indicate a duplicate MEP or network loop. If these error types are not found, the CCM is stored in the MEP's local database until aged out.
- ◆ If a maintenance point fails to receive three consecutive CCMs from any other MEP in the same MA, a connectivity failure is registered.
- ◆ If a maintenance point receives a CCM with an invalid MEPID or MA level or an MA level lower than its own, a failure is registered which indicates a configuration error or cross-connect error (i.e., overlapping MAs).
- ◆ The interval at which CCMs are issued should be configured to detect connectivity problems in a timely manner, as dictated by the nature and size of the MA.
- ◆ The maintenance of a MIP CCM database by a MIP presents some difficulty for bridges carrying a large number of Service Instances, and for whose MEPs are issuing CCMs at a high frequency. For this reason, slower CCM transmission rates may have to be used.

Parameters

These parameters are displayed:

Creating a Maintenance Association

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **MA Name** – MA name. (Range: 1-43¹³ alphanumeric characters)
Each MA name must be unique within the CFM domain.
- ◆ **Primary VLAN** – Service VLAN ID. (Range: 1-4094)
This is the VLAN through which all CFM functions are executed for this MA.
- ◆ **MIP Creation Type** – Specifies the CFM protocol's creation method for maintenance intermediate points (MIPs) in this MA:
 - **Default** – MIPs can be created for this MA on any bridge port through which the MA's VID can pass.
 - **Explicit** – MIPs can be created for this MA only on bridge ports through which the MA's VID can pass, and only if a maintenance end point (MEP) is created at some lower MA Level.
 - **None** – No MIP can be created for this MA.

Configuring Detailed Settings for a Maintenance Association

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **MA Name Format** – Specifies the name format for the maintenance association as IEEE 802.1ag character based, or ITU-T SG13/SG15 Y.1731 defined ICC-based format.
 - **Character String** – IEEE 802.1ag defined character string format. This is an IETF RFC 2579 DisplayString.
 - **ICC Based** – ITU-T SG13/SG15 Y.1731 defined ICC based format.
- ◆ **Interval Level** – The delay between sending CCMs. The setting for this parameter is expressed as levels 4 through 7, which in turn map to specific intervals of time. (Options: 4 - 1 second, 5 - 10 seconds, 6 - 1 minute, 7 - 10 minutes)
- ◆ **Connectivity Check** – Enables transmission of CCMs. (Default: Disabled)
- ◆ **Cross Check** – Enables cross-checking between a static list of MEPs assigned to other devices within the same maintenance association and the MEPs learned through CCMs.

13. The total length of the MD name and MA name cannot exceed 44 characters.

Before starting the cross-check process, first configure the remote MEPs that exist on other devices inside the maintenance association using the Remote MEP List (see "[Configuring Remote Maintenance End Points](#)"). These remote MEPs are used in the cross-check operation to verify that all endpoints in the specified MA are operational.

The cross-check start delay, which sets the maximum delay this device waits for a remote MEP to come up before starting the cross-check operation, is a domain-level parameter. To set this parameter, use the CFM MD Configuration screen (see "[Configuring CFM Maintenance Domains](#)").

- ◆ **AIS Status** – Enables/disables suppression of the Alarm Indication Signal (AIS). (Default: Disabled)
- ◆ **AIS Period** – Configures the period at which AIS is sent in an MA. (Range: 1 or 60 seconds; Default: 1 second)
- ◆ **AIS Transmit Level** – Configure the AIS maintenance level in an MA. (Range: 0-7; Default is 0)
AIS Level must follow this rule: AIS Level \geq Domain Level
- ◆ **AIS Suppress Alarm** – Enables/disables suppression of the AIS. (Default: Disabled)

Web Interface

To create a maintenance association:

1. Click Administration, CFM.
2. Select Configure MA from the Step list.
3. Select Add from the Action list.
4. Select an entry from the MD Index list.
5. Specify the MAs assigned to each domain, the VLAN through which CFM messages are passed, and the manner in which MIPs can be created within each MA.
6. Click Apply.

Figure 249: Creating Maintenance Associations

The screenshot shows the 'Administration > CFM' configuration page. At the top, there is a breadcrumb trail 'Administration > CFM'. Below it, there are two dropdown menus: 'Step: 4. Configure MA' and 'Action: Add'. The main form contains the following fields:

- MD Index: A dropdown menu with '1' selected.
- MA Index (1-2147483647): A text input field containing '1'.
- MA Name: A text input field containing 'rd'.
- Primary VLAN (1-4094): A text input field containing '1'.
- MIP Creation Type: A dropdown menu with 'Default' selected.

At the bottom right of the form, there are two buttons: 'Apply' and 'Revert'.

To show the configured maintenance associations:

1. Click Administration, CFM.
2. Select Configure MA from the Step list.
3. Select Show from the Action list.
4. Select an entry from the MD Index list.

Figure 250: Showing Maintenance Associations

The screenshot shows the 'Administration > CFM' configuration page with the 'Action' dropdown set to 'Show'. The 'MD Index' dropdown is set to '1'. Below this, there is a table titled 'CFM MA List' with a 'Total: 1' indicator. The table has the following structure:

<input type="checkbox"/>	MA Index	MA Name	Primary VLAN	MIP Creation Type
<input type="checkbox"/>	1	rd	1	Default

At the bottom right of the table, there are two buttons: 'Delete' and 'Revert'.

To configure detailed settings for maintenance associations:

1. Click Administration, CFM.
2. Select Configure MA from the Step list.
3. Select Configure Details from the Action list.
4. Select an entry from MD Index and MA Index.
5. Specify the CCM interval, enable the transmission of connectivity check and cross check messages, and configure the required AIS parameters.
6. Click Apply

Figure 251: Configuring Detailed Settings for Maintenance Associations

The screenshot shows the 'Administration > CFM' configuration page. At the top, there are two dropdown menus: 'Step: 4. Configure MA' and 'Action: Configure Details'. Below this, the configuration parameters are listed as follows:

MD Index	1
MA Index	1
MA Name Format	Character String
Interval Level (4-7)	4
Connectivity Check	<input checked="" type="checkbox"/> Enabled
Cross Check	<input checked="" type="checkbox"/> Enabled
AIS Status	<input checked="" type="checkbox"/> Enabled
AIS Period	1
AIS Transmit Level (0-7)	0
AIS Suppress Alarm	<input type="checkbox"/> Enabled

At the bottom right of the configuration area, there are two buttons: 'Apply' and 'Revert'.

Configuring Maintenance End Points

Use the Administration > CFM (Configure MEP – Add) page to configure Maintenance End Points (MEPs). MEPs, also called Domain Service Access Points (DSAPs), must be configured at the domain boundary to provide management access for each maintenance association.

Command Usage

- ◆ CFM elements must be configured in the following order: (1) maintenance domain at the same level as the MEP to be configured (see "[Configuring CFM Maintenance Domains](#)"), (2) maintenance association within the domain (see "[Configuring CFM Maintenance Associations](#)"), and (3) finally the MEPs using the MEP List.
- ◆ An interface may belong to more than one domain, or to different MAs in different domains.
- ◆ To change the MEP's MA or the direction it faces, first delete the MEP, and then create a new one.

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **MEP ID** – Maintenance end point identifier. (Range: 1-8191)
- ◆ **MEP Direction** – Up indicates that the MEP faces inward toward the switch cross-connect matrix, and transmits CFM messages towards, and receives them from, the direction of the internal bridge relay mechanism. If the **Up** option is

not selected, then the MEP is facing away from the switch, and transmits CFM messages towards, and receives them from, the direction of the physical medium.

- ◆ **Interface** – Indicates a port or trunk.

Web Interface

To configure a maintenance end point:

1. Click Administration, CFM.
2. Select Configure MEP from the Step list.
3. Select Add from the Action list.
4. Select an entry from MD Index and MA Index.
5. Specify the MEPs assigned to each MA, set the MEP identifier, the direction in which the MEP faces, and the physical interface serving as the DSAP.
6. Click Apply.

Figure 252: Configuring Maintenance End Points

Administration > CFM

Step: 3. Configure MEP Action: Add

MD Index 1

MA Index 1

MEP ID (1-8191) 1

MEP Direction Up

Interface Port 1 Trunk

Apply Revert

To show the configured maintenance end points:

1. Click Administration, CFM.
2. Select Configure MEP from the Step list.
3. Select Show from the Action list.

4. Select an entry from MD Index and MA Index.

Figure 253: Showing Maintenance End Points



Configuring Remote Maintenance End Points

Use the Administration > CFM (Configure Remote MEP – Add) page to specify remote maintenance end points (MEPs) set on other CFM-enabled devices within a common MA. Remote MEPs can be added to a static list in this manner to verify that each entry has been properly configured and is operational. When cross-checking is enabled, the list of statically configured remote MEPs is compared against the MEPs learned through continuity check messages (CCMs), and any discrepancies reported via SNMP traps.

Command Usage

- ◆ All MEPs that exist on other devices inside a maintenance association should be statically configured to ensure full connectivity through the cross-check process.
- ◆ Remote MEPs can only be configured if local domain service access points (DSAPs) have already been created (see "[Configuring Maintenance End Points](#)") at the same maintenance level and in the same MA. DSAPs are MEPs that exist on the edge of the domain, and act as primary service access points for end-to-end cross-check, loop-back, and link-trace functions.
- ◆ The MEP cross-check start delay which sets the maximum delay that a device waits for remote MEPs to come up before starting the cross-check operation can be configured on the Configure Global page (see "[Configuring Global Settings for CFM](#)").
- ◆ SNMP traps for continuity check events discovered by cross-check operations can also be configured on the Configure Global page (see "[Configuring Global Settings for CFM](#)").

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)

- ◆ **MEP ID** – Identifier for a maintenance end point which exists on another CFM-enabled device within the same MA. (Range: 1-8191)

Web Interface

To configure a remote maintenance end point:

1. Click Administration, CFM.
2. Select Configure Remote MEP from the Step list.
3. Select Add from the Action list.
4. Select an entry from MD Index and MA Index.
5. Specify the remote MEPs which exist on other devices within the same MA.
6. Click Apply.

Figure 254: Configuring Remote Maintenance End Points

The screenshot shows the 'Administration > CFM' page. The 'Step' dropdown is set to '4. Configure Remote MEP' and the 'Action' dropdown is set to 'Add'. There are three input fields: 'MD Index' with a dropdown menu showing '1', 'MA Index' with a dropdown menu showing '1', and 'MEP ID (1-8191)' with a text input field containing '2'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show the configured remote maintenance end points:

1. Click Administration, CFM.
2. Select Configure MEP from the Step list.
3. Select Show from the Action list.
4. Select an entry from MD Index and MA Index.

Figure 255: Showing Remote Maintenance End Points

The screenshot shows the 'Administration > CFM' page. The 'Step' dropdown is set to '6. Configure Remote MEP' and the 'Action' dropdown is set to 'Show'. There are two input fields: 'MD Index' with a dropdown menu showing '1' and 'MA Index' with a dropdown menu showing '1'. Below these fields, there is a table titled 'CFM Remote Maintenance Association End Point List' with a 'Total: 1' indicator. The table has two columns: a checkbox column and a 'MEP ID' column. The first row has a checked checkbox and the value '2' in the 'MEP ID' column. At the bottom right, there are 'Delete' and 'Revert' buttons.

	MEP ID
<input checked="" type="checkbox"/>	2

Transmitting Link Trace Messages Use the Administration > CFM (Transmit Link Trace) page to transmit link trace messages (LTMs). These messages can isolate connectivity faults by tracing the path through a network to the designated target node (i.e., a remote maintenance end point).

Command Usage

- ◆ LTMs can be targeted to MEPs, not MIPs. Before sending a link trace message, be sure you have configured the target MEP for the specified MA (see ["Configuring Remote Maintenance End Points"](#)).
- ◆ If MAC address of target MEP has not been learned by any local MEP, then the linktrace may fail. Use the Show Remote MEP page (see ["Displaying Remote MEPs"](#)) to verify that a MAC address has been learned for the target MEP.
- ◆ LTMs are sent as multicast CFM frames, and forwarded from MIP to MIP, with each MIP generating a link trace reply, up to the point at which the LTM reaches its destination or can no longer be forwarded.
- ◆ LTMs are used to isolate faults. However, this task can be difficult in an Ethernet environment, since each node is connected through multipoint links. Fault isolation is even more challenging since the MAC address of the target node can age out in several minutes. This can cause the traced path to vary over time, or connectivity lost if faults cause the target MEP to be isolated from other MEPs in an MA.
- ◆ When using the command line or web interface, the source MEP used by to send a link trace message is chosen by the CFM protocol. However, when using SNMP, the source MEP can be specified by the user.
- ◆ Parameters controlling the link trace cache, including operational state, entry hold time, and maximum size can be configured on the Configure Global page (see ["Configuring Global Settings for CFM"](#)).

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **Source MEP ID** – The identifier of a source MEP that will send the link trace message. (Range: 1-8191)
- ◆ **Target**
 - **MEP ID** – The identifier of a remote MEP that is the target of a link trace message. (Range: 1-8191)
 - **MAC Address** – MAC address of a remote MEP that is the target of a link trace message. This address can be entered in either of the following formats: xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx

- ◆ **TTL** – The time to live of the link trace message. (Range: 0-255 hops)

Web Interface

To transmit link trace messages:

1. Click Administration, CFM.
2. Select Transmit Link Trace from the Step list.
3. Select an entry from MD Index and MA Index.
4. Specify the source MEP, the target MEP using either its MEP identifier or MAC address, and set the maximum number of hops allowed in the TTL field.
5. Click Apply.
6. Check the results in the Link Trace cache (see "[Displaying the Link Trace Cache](#)").

Figure 256: Transmitting Link Trace Messages

The screenshot shows a web interface titled "Administration > CFM". At the top, there is a "Step:" dropdown menu set to "7. Transmit Link Trace". Below this, there are several input fields: "MD Index" with a dropdown menu showing "1", "MA Index" with a dropdown menu showing "1", "Source MEP ID (1-8191)" with a text input field containing "1", "Target" with a radio button selected for "MEP ID (1-8191)" and a text input field containing "2", and "TTL (0-255)" with a text input field containing "5". There is also an unselected radio button for "MAC Address" with an empty text input field. At the bottom right, there are two buttons: "Apply" and "Revert".

Transmitting Loop Back Messages

Use the Administration > CFM (Transmit Loopback) page to transmit Loopback Messages (LBMs). These messages can be used to isolate or verify connectivity faults by submitting a request to a target node (i.e., a remote MEP or MIP) to echo the message back to the source.

Command Usage

- ◆ Loopback messages can be used for fault verification and isolation after automatic detection of a fault or receipt of some other error report. Loopback messages can also be used to confirm the successful restoration or initiation of connectivity. The receiving maintenance point should respond to the loop back message with a loopback reply.
- ◆ The point from which the loopback message is transmitted (i.e., a local DSAP) and the target maintenance point must be within the same MA.

- ◆ If the continuity check database does not have an entry for the specified maintenance point, an error message will be displayed.
- ◆ When using the command line or web interface, the source MEP used by to send a loopback message is chosen by the CFM protocol. However, when using SNMP, the source MEP can be specified by the user.

Parameters

These parameters are displayed:

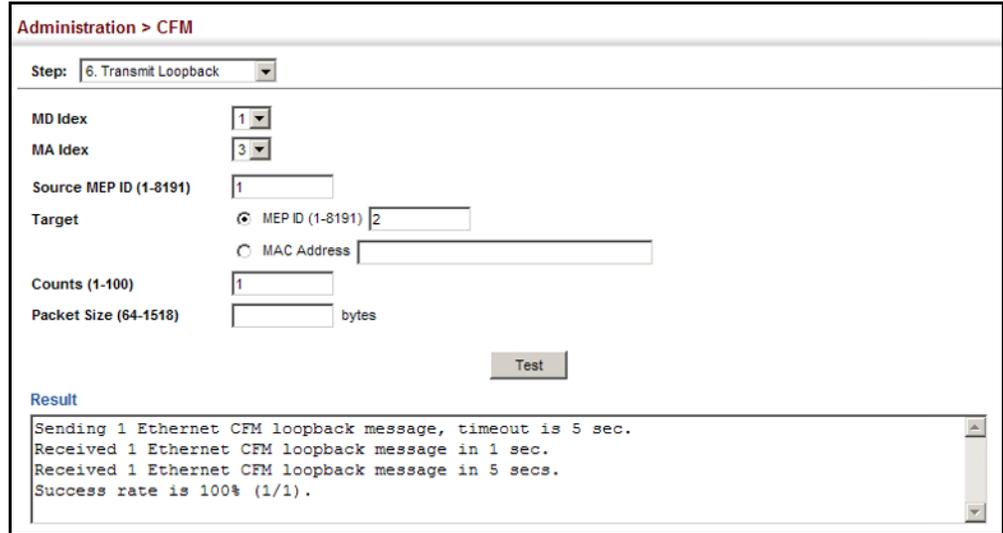
- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **Source MEP ID** – The identifier of a source MEP that will send the loopback message. (Range: 1-8191)
- ◆ **Target**
 - **MEP ID** – The identifier of a remote MEP that is the target of a loopback message. (Range: 1-8191)
 - **MAC Address** – MAC address of a remote MEP that is the target of a loopback message. This address can be entered in either of the following formats: xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx
- ◆ **Count** – The number of times the loopback message is sent. (Range: 1-1024)
- ◆ **Packet Size** – The size of the loopback message. (Range: 64-1518 bytes; Default: 64 bytes)

Web Interface

To transmit loopback messages:

1. Click Administration, CFM.
2. Select Transmit Loopback from the Step list.
3. Select an entry from MD Index and MA Index.
4. Specify the source MEP, the target MEP using either its MEP identifier or MAC address, set the number of times the loopback message is to be sent.
5. Click Apply.

Figure 257: Transmitting Loopback Messages



Transmitting Delay-Measure Requests Use the Administration > CFM (Transmit Delay Measure) page to send periodic delay-measure requests to a specified MEP within a maintenance association.

Command Usage

- ◆ Delay measurement can be used to measure frame delay and frame delay variation between MEPs.
- ◆ A local MEP must be configured for the same MA before you can use this function.
- ◆ If a MEP is enabled to generate frames with delay measurement (DM) information, it periodically sends DM frames to its peer MEP in the same MA., and expects to receive DM frames back from it.

- ◆ Frame delay measurement can be made only for two-way measurements, where the MEP transmits a frame with DM request information with the TxTimeStampf (Timestamp at the time of sending a frame with DM request information), and the receiving MEP responds with a frame with DM reply information with TxTimeStampf copied from the DM request information, RxTimeStampf (Timestamp at the time of receiving a frame with DM request information), and TxTimeStamptb (Timestamp at the time of transmitting a frame with DM reply information):

$$\text{Frame Delay} = (\text{RxTimeStampb} - \text{TxTimeStampf}) - (\text{TxTimeStamptb} - \text{RxTimeStampf})$$

- ◆ The MEP can also make two-way frame delay variation measurements based on its ability to calculate the difference between two subsequent two-way frame delay measurements.

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **Source MEP ID** – The identifier of a source MEP that will send the delay-measure message. (Range: 1-8191)
- ◆ **Target**
 - **MEP ID** – The identifier of a remote MEP that is the target of a delay-measure message. (Range: 1-8191)
 - **MAC Address** – MAC address of a remote MEP that is the target of a delay-measure message. This address can be entered in either of the following formats: xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx
- ◆ **Count** – The number of times to retry sending the message if no response is received before the specified timeout. (Range: 1-5; Default: 5)
- ◆ **Packet Size** – The size of the delay-measure message. (Range: 64-1518 bytes; Default: 64 bytes)
- ◆ **Interval** – The transmission delay between delay-measure messages. (Range: 1-5 seconds; Default: 1 second)
- ◆ **Timeout** – The timeout to wait for a response. (Range: 1-5 seconds; Default: 5 seconds)

Web Interface

To transmit delay-measure messages:

1. Click Administration, CFM.
2. Select Transmit Delay Measure from the Step list.
3. Select an entry from MD Index and MA Index.
4. Specify the source MEP, the target MEP using either its MEP identifier or MAC address, set the number of times the delay-measure message is to be sent, the interval, and the timeout.
5. Click Apply.

Figure 258: Transmitting Delay-Measure Messages

Administration > CFM

Step: 7. Transmit Delay Measure

MD Index: 1

MA Index: 3

Source MEP ID (1-8191):

Target: MEP ID (1-8191) MAC Address

Counts (1-5):

Packet Size (64-1518): bytes

Interval (1-5): ms

Timeout (1-5): ms

Apply Revert

Result

```
Sending 5 Ethernet CFM delay measurement message, timeout is 5 sec.
Sequence  Delay Time (ms.)  Delay Variation (ms.)
-----
1          < 10              0
2          10              10
3          < 10              10
4          < 10              0
5          < 10              0
Success rate is 100% (5/5), delay time min/avg/max=0/2/10 ms.
Average frame delay variation is 4 ms.
```

Displaying Local MEPs Use the Administration > CFM > Show Information (Show Local MEP) page to show information for the MEPs configured on this device.

Parameters

These parameters are displayed:

- ◆ **MEP ID** – Maintenance end point identifier.
- ◆ **MD Name** – Maintenance domain name.
- ◆ **Level** – Authorized maintenance level for this domain.
- ◆ **Direction** – Direction in which the MEP communicates CFM messages:
 - Down indicates that the MEP is facing away from the switch, and transmits CFM messages towards, and receives them from, the direction of the physical medium.
 - Up indicates that the MEP faces inward toward the switch cross-connect matrix, and transmits CFM messages towards, and receives them from, the direction of the internal bridge relay mechanism.
- ◆ **Primary VLAN** – Service VLAN ID.
- ◆ **Interface** – Physical interface of this entry (either a port or trunk).
- ◆ **CC Status** – Shows administrative status of CCMs.

- ◆ **MAC Address** – MAC address of this MEP entry.

Web Interface

To show information for the MEPs configured on this device:

1. Click Administration, CFM.
2. Select Show Information from the Step list.
3. Select Show Local MEP from the Action list.

Figure 259: Showing Information on Local MEPs

Administration > CFM							
Step: 10. Show Information		Action: Show Local MEP					
CFM Local Maintenance Association End Point Information Total: 1							
MEP ID	MD Name	Level	Direction	Primary VLAN	Interface	CC Status	MAC Address
1	voip	3	Down	1	Unit 1 / Port 2	Enabled	00-E0-0C-00-00-FF

Displaying Details for Local MEPs

Use the Administration > CFM > Show Information (Show Local MEP Details) page to show detailed CFM information about a local MEP in the continuity check database.

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **MEP ID** – Maintenance end point identifier. (Range: 1-8191)
- ◆ **MD Name** – The maintenance domain for this entry.
- ◆ **MA Name** – Maintenance association to which this remote MEP belongs.
- ◆ **MA Name Format** – The format of the Maintenance Association name, including primary VID, character string, unsigned Integer 16, or RFC 2865 VPN ID.
- ◆ **Level** – Maintenance level of the local maintenance point.
- ◆ **Direction** – The direction in which the MEP faces on the Bridge port (up or down).
- ◆ **Interface** – The port to which this MEP is attached.
- ◆ **CC Status** – Shows if the MEP will generate CCM messages.

- ◆ **MAC Address** – MAC address of the local maintenance point. (If a CCM for the specified remote MEP has never been received or the local MEP record times out, the address will be set to the initial value of all Fs.)
- ◆ **Defect Condition** – Shows the defect detected on the MEP.
- ◆ **Received RDI** – Receive status of remote defect indication (RDI) messages on the MEP.
- ◆ **AIS Status** – Shows if MEPs within the specified MA are enabled to send frames with AIS information following detection of defect conditions.
- ◆ **AIS Period** – The interval at which AIS information is sent.
- ◆ **AIS Transmit Level** – The maintenance level at which AIS information will be sent for the specified MEP.
- ◆ **Suppress Alarm** – Shows if the specified MEP is configured to suppress sending frames containing AIS information following the detection of defect conditions.
- ◆ **Suppressing Alarms** – Shows if the specified MEP is currently suppressing sending frames containing AIS information following the detection of defect conditions.

Web Interface

To show detailed information for the MEPs configured on this device:

1. Click Administration, CFM.
2. Select Show Information from the Step list.
3. Select Show Local MEP Details from the Action list.
4. Select an entry from MD Index and MA Index.
5. Select a MEP ID.

Figure 260: Showing Detailed Information on Local MEPs

Administration > CFM

Step: 10. Show Information Action: Show Local MEP Details

MD Index: 1
 MA Index: 1
 MEP ID: 1

Query

MD Name	md1
MA Name	ma1
MA Name Format	Character String
Level	0
Direction	Up
Interface	Unit 1 / Port 2
CC Status	Enabled
MAC Address	00-1B-D5-50-91-FD
Defect Condition	No Defect
Received RDI	False
AIS Status	Enabled
AIS Period	60 sec
AIS Transmit Level	Default
Suppress Alarm	Enabled
Suppressing Alarms	Disabled

Displaying Local MIPs Use the Administration > CFM > Show Information (Show Local MIP) page to show the MIPs on this device discovered by the CFM protocol. (For a description of MIPs, refer to the Command Usage section under "[Configuring CFM Maintenance Domains](#)".)

Parameters

These parameters are displayed:

- ◆ **MD Name** – Maintenance domain name.
- ◆ **Level** – Authorized maintenance level for this domain.
- ◆ **MA Name** – Maintenance association name.
- ◆ **Primary VLAN** – Service VLAN ID.
- ◆ **Interface** – Physical interface of this entry (either a port or trunk).

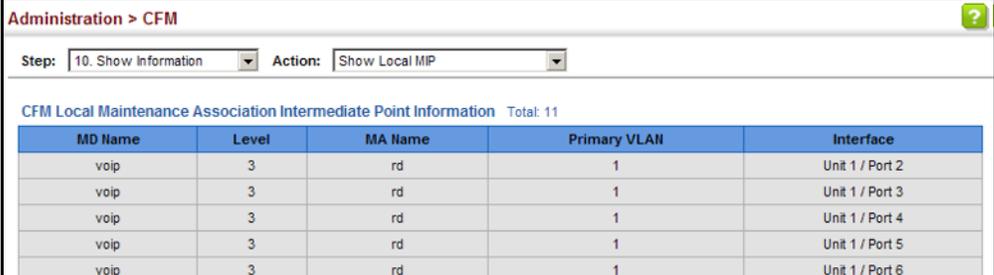
Web Interface

To show information for the MIPs discovered by the CFM protocol:

1. Click Administration, CFM.

2. Select Show Information from the Step list.
3. Select Show Local MIP from the Action list.

Figure 261: Showing Information on Local MIPs



The screenshot shows the Administration > CFM web interface. At the top, there is a breadcrumb trail 'Administration > CFM' and a help icon. Below this, there are two dropdown menus: 'Step: 10. Show Information' and 'Action: Show Local MIP'. The main content area displays a table titled 'CFM Local Maintenance Association Intermediate Point Information' with a total of 11 entries. The table has five columns: MD Name, Level, MA Name, Primary VLAN, and Interface. The data rows show five entries for MD Name 'voip', Level '3', MA Name 'rd', Primary VLAN '1', and Interface 'Unit 1 / Port 2' through 'Unit 1 / Port 6'.

MD Name	Level	MA Name	Primary VLAN	Interface
voip	3	rd	1	Unit 1 / Port 2
voip	3	rd	1	Unit 1 / Port 3
voip	3	rd	1	Unit 1 / Port 4
voip	3	rd	1	Unit 1 / Port 5
voip	3	rd	1	Unit 1 / Port 6

Displaying Remote MEPs

Use the Administration > CFM > Show Information (Show Remote MEP) page to show MEPs located on other devices which have been discovered through continuity check messages, or statically configured in the MEP database and verified through cross-check messages.

Parameters

These parameters are displayed:

- ◆ **MEP ID** – Maintenance end point identifier.
- ◆ **MA Name** – Maintenance association name.
- ◆ **Level** – Authorized maintenance level for this domain.
- ◆ **Primary VLAN** – Service VLAN ID.
- ◆ **MEP Up** – Indicates whether or not this MEP is functioning normally.
- ◆ **Remote MAC Address** – MAC address of the remote maintenance point. (If a CCM for the specified remote MEP has never been received or the remote MEP record times out, the address will be set to the initial value of all Fs.)

Web Interface

To show information for remote MEPs:

1. Click Administration, CFM.
2. Select Show Information from the Step list.
3. Select Show Remote MEP from the Action list.

Figure 262: Showing Information on Remote MEPs

Administration > CFM

Step: 10. Show Information Action: Show Remote MEP

CFM Remote Maintenance Association End Point Information Total: 1

MEP ID	MA Name	Level	Primary VLAN	MEP Up	Remote MAC Address
2	rd	3	1	Yes	74-8E-F8-68-02-32

Clear

Displaying Details for Remote MEPs

Use the Administration > CFM > Show Information (Show Remote MEP Details) page to show detailed information for MEPs located on other devices which have been discovered through continuity check messages, or statically configured in the MEP database and verified through cross-check messages.

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **MEP ID** – Maintenance end point identifier. (Range: 1-8191)
- ◆ **MD Name** – Maintenance domain name.
- ◆ **MA Name** – Maintenance association name.
- ◆ **Level** – Authorized maintenance level for this domain.
- ◆ **MAC Address** – MAC address of this MEP entry.
- ◆ **Primary VLAN** – Service VLAN ID.
- ◆ **Incoming Port** – Port to which this remote MEP is attached.
- ◆ **CC Lifetime** – Length of time to hold messages about this MEP in the CCM database.
- ◆ **Age of Last CC Message** – Length of time the last CCM message about this MEP has been in the CCM database.
- ◆ **Frame Loss** – Percentage of transmitted frames lost.
- ◆ **CC Packet Statistics** – The number of CCM packets received successfully and those with errors.
- ◆ **Port State** – Port states include:
 - Up – The port is functioning normally.

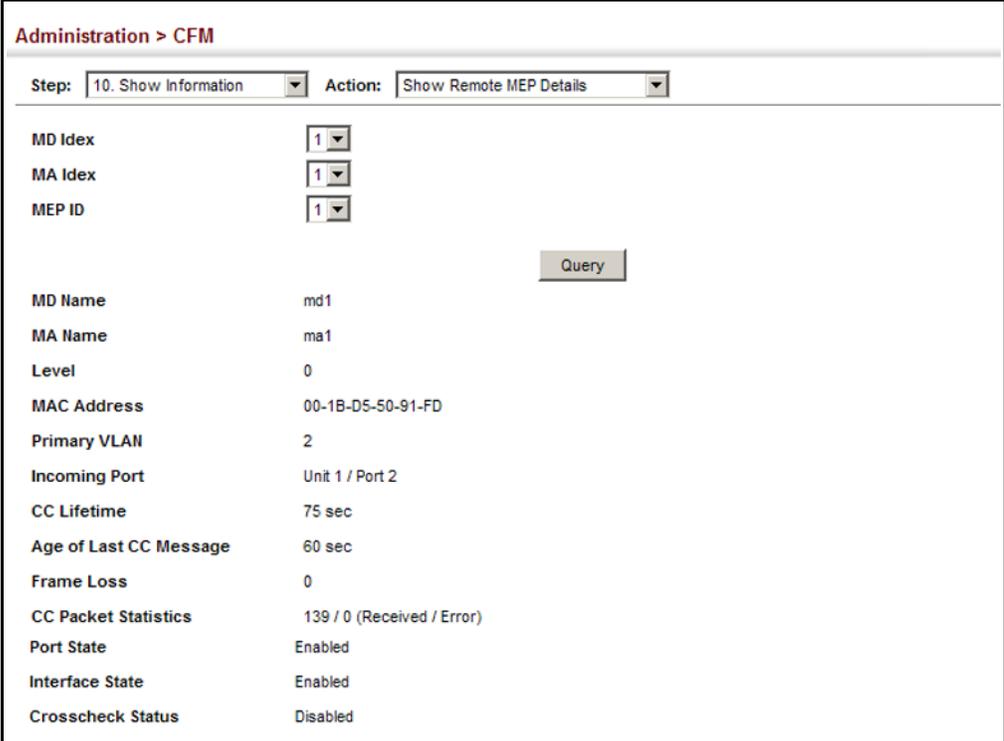
- Blocked – The port has been blocked by the Spanning Tree Protocol.
- No port state – Either no CCM has been received, or no port status TLV was received in the last CCM.
- ◆ **Interface State** – Interface states include:
 - No Status – Either no CCM has been received, or no interface status TLV was received in the last CCM.
 - Up – The interface is ready to pass packets.
 - Down – The interface cannot pass packets.
 - Testing – The interface is in some test mode.
 - Unknown – The interface status cannot be determined for some reason.
 - Dormant – The interface is not in a state to pass packets but is in a pending state, waiting for some external event.
 - Not Present – Some component of the interface is missing.
 - isLowerLayerDown – The interface is down due to state of the lower layer interfaces.
- ◆ **Crosscheck Status** – Shows if crosscheck function has been enabled.

Web Interface

To show detailed information for remote MEPs:

1. Click Administration, CFM.
2. Select Show Information from the Step list.
3. Select Show Remote MEP Details from the Action list.
4. Select an entry from MD Index and MA Index.
5. Select a MEP ID.

Figure 263: Showing Detailed Information on Remote MEPs



Administration > CFM

Step: 10. Show Information Action: Show Remote MEP Details

MD Index 1
MA Index 1
MEP ID 1

Query

MD Name	md1
MA Name	ma1
Level	0
MAC Address	00-1B-D5-50-91-FD
Primary VLAN	2
Incoming Port	Unit 1 / Port 2
CC Lifetime	75 sec
Age of Last CC Message	60 sec
Frame Loss	0
CC Packet Statistics	139 / 0 (Received / Error)
Port State	Enabled
Interface State	Enabled
Crosscheck Status	Disabled

Displaying the Link Trace Cache Use the Administration > CFM > Show Information (Show Link Trace Cache) page to show information about link trace operations launched from this device.

Parameters

These parameters are displayed:

- ◆ **Hops** – The number hops taken to reach the target MEP.
- ◆ **MA** – Maintenance association name.
- ◆ **IP/Alias** – IP address or DNS alias of the target device’s CPU.
- ◆ **Forwarded** – Shows whether or not this link trace message was forwarded. A message is not forwarded if received by the target MEP.
- ◆ **Ingress MAC Address** – MAC address of the ingress port on the target device.
- ◆ **Egress MAC Address** – MAC address of the egress port on the target device.
- ◆ **Ingress Action** – Action taken on the ingress port:
 - IngOk – The target data frame passed through to the MAC Relay Entity.
 - IngDown – The bridge port’s MAC_Operational parameter is false. This value could be returned, for example, by an operationally Down MEP that

has another Down MEP at a higher MD level on the same bridge port that is causing the bridge port's MAC_Operational parameter to be false.

- IngBlocked – The ingress port can be identified, but the target data frame was not forwarded when received on this port due to active topology management, i.e., the bridge port is not in the forwarding state.
- IngVid – The ingress port is not in the member set of the LTM's VIDs, and ingress filtering is enabled, so the target data frame was filtered by ingress filtering.

◆ **Egress Action** – Action taken on the egress port:

- EgrOk – The targeted data frame was forwarded.
- EgrDown – The Egress Port can be identified, but that bridge port's MAC_Operational parameter is false.
- EgrBlocked – The egress port can be identified, but the data frame was not passed through the egress port due to active topology management, i.e., the bridge port is not in the forwarding state.
- EgrVid – The Egress Port can be identified, but the bridge port is not in the LTM's VID member set, and was therefore filtered by egress filtering.

◆ **Reply** – Reply action:

- FDB – Target address found in forwarding database.
- MPDB – Target address found in the maintenance point database.
- HIT – Target located on this device.

Web Interface

To show information about link trace operations launched from this device:

1. Click Administration, CFM.
2. Select Show Information from the Step list.
3. Select Show Link Trace Cache from the Action list.

Figure 264: Showing the Link Trace Cache

Hops	MA	IP/Alias Address	Forwarded	Ingress MAC Address	Egress MAC Address	Ingress Action	Egress Action	Reply
1	rd	192.168.0.3	Not Forwarded	74-8E-F8-68-02-32		ingOk		Hit

Displaying Fault Notification Settings

Use the Administration > CFM > Show Information (Show Fault Notification Generator) page to display configuration settings for the fault notification generator.

Parameters

These parameters are displayed:

- ◆ **MEP ID** – Maintenance end point identifier.
- ◆ **MD Name** – Maintenance domain name.
- ◆ **MA Name** – Maintenance association name.
- ◆ **Highest Defect** – The highest defect that will generate a fault alarm. (This is disabled by default.)
- ◆ **Lowest Alarm** – The lowest defect that will generate a fault alarm¹⁴.
- ◆ **Alarm Time** – The time a defect must exist before a fault alarm is issued¹⁴.
- ◆ **Reset Time** – The time after a fault alarm has been issued, and no defect exists, before another fault alarm can be issued¹⁴.

Web Interface

To show configuration settings for the fault notification generator:

1. Click Administration, CFM.
2. Select Show Information from the Step list.
3. Select Show Fault Notification Generator from the Action list.

Figure 265: Showing Settings for the Fault Notification Generator

Administration > CFM						
Step: 10. Show Information		Action: Show Fault Notification Generator				
CFM Fault Notification Generator Information Total: 1						
MEP ID	MD Name	MA Name	Highest Defect	Lowest Alarm	Alarm Time (sec)	Reset Time (sec)
1	voip	rd	defRemoteCCM	macRemErrXcon	3	10

Displaying Continuity Check Errors

Use the Administration > CFM > Show Information (Show Continuity Check Error) page to display the CFM continuity check errors logged on this device.

Parameters

These parameters are displayed:

- ◆ **Level** – Maintenance level associated with this entry.

¹⁴. See “Configuring CFM Maintenance Domains” on page 398.

- ◆ **Primary VLAN** – VLAN in which this error occurred.
- ◆ **MEP ID** – Identifier of remote MEP.
- ◆ **Interface** – Port at which the error was recorded.
- ◆ **Remote MAC** – MAC address of remote MEP.
- ◆ **Reason** – Error types include:
 - LEAK – MA x is associated with a specific VID list¹⁵, one or more of the VIDs in this MA can pass through the bridge port, no MEP is configured facing outward (down) on any bridge port for this MA, and some other MA y, at a higher maintenance level, and associated with at least one of the VID(s) also in MA x, does have a MEP configured on the bridge port.
 - VIDS – MA x is associated with a specific VID list¹⁵, an MEP is configured facing inward (up) on this MA on the bridge port, and some other MA y, associated with at least one of the VID(s) also in MA x, also has an Up MEP configured facing inward (up) on some bridge port.
 - EXCESS_LEV – The number of different MD levels at which MIPs are to be created on this port exceeds the bridge's capabilities.
 - OVERLAP_LEV – A MEP is created for one VID at one maintenance level, but a MEP is configured on another VID at an equivalent or higher level, exceeding the bridge's capabilities.
- ◆ **MA Name** – The maintenance association for this entry.

Web Interface

To show CFM continuity check errors:

1. Click Administration, CFM.
2. Select Show Information from the Step list.
3. Select Show Continuity Check Error from the Action list.

Figure 266: Showing Continuity Check Errors

Administration > CFM						
Step: 10. Show Information		Action: Show Continuity Check Error				
CFM Continuity Check Error Information Total: 1						
Level	Primary VLAN	MEP ID	Interface	Remote MAC	Reason	MA Name
5	2	40	Unit 1 / Port 10	00-01-02-03-04-05	LEAK	aa

Clear

15. This definition is based on the IEEE 802.1ag standard. Current software for this switch only supports a single VLAN per MA. However, since it may interact with other devices which support multiple VLAN assignments per MA, this error message may be reported.

UDLD Configuration

The switch can be configured to detect general loopback conditions caused by hardware problems or faulty protocol settings. When enabled, a control frame is transmitted on the participating ports, and the switch monitors inbound traffic to see if the frame is looped back.

Usage Guidelines

- ◆ The default settings for the control frame transmit interval and recover time may be adjusted to improve performance for your specific environment. The shutdown mode may also need to be changed once you determine what kind of packets are being looped back.
- ◆ General loopback detection provided by the commands described in this section and loopback detection provided by the spanning tree protocol cannot both be enabled at the same time. If loopback detection is enabled for the spanning tree protocol, general loopback detection cannot be enabled on the same interface.
- ◆ When a loopback event is detected on an interface or when an interface is released from a shutdown state caused by a loopback event, a trap message is sent and the event recorded in the system log.
- ◆ Loopback detection must be enabled both globally and on an interface for loopback detection to take effect.

Configuring UDLD Protocol Intervals

Use the Administration > UDLD > Configure Global page to configure the UniDirectional Link Detection message probe interval, detection interval, and recovery interval.

Parameters

These parameters are displayed:

- ◆ **Message Interval** – Configures the message interval between UDLD probe messages for ports in the advertisement phase and determined to be bidirectional. (Range: 7-90 seconds; Default: 15 seconds)

UDLD probe messages are sent after linkup or detection phases. During the detection phase, messages are exchanged at the maximum rate of one per second. After that, if the protocol reaches a stable state and determines that the link is bidirectional, the message interval is increased to a configurable value based on a curve known as $M1(t)$, a time-based function described in RFC 5171.

If the link is deemed anything other than bidirectional at the end of the detection phase, this curve becomes a flat line with a fixed value of M_{fast} (7 seconds).

If the link is instead deemed bidirectional, the curve will use M_{fast} for the first four subsequent message transmissions and then transition to an M_{slow} value

for all other steady-state transmissions. Mslow is the value configured by this command.

- ◆ **Detection Interval** – Sets the amount of time the switch remains in detection state after discovering a neighbor. (Range: 5-255 seconds; Default: 5 seconds)

When a neighbor device is discovered by UDLD, the switch enters “detection state” and remains in this state for specified detection-interval. After the detection-interval expires, the switch tries to decide whether or the link is unidirectional based on the information collected during the “detection state.”

- ◆ **Recovery Status** – Configures the switch to automatically recover from UDLD disabled port state after a period specified by the Recovery Interval. (Default: Disabled)

When automatic recovery state is changed, any ports shut down by UDLD will be reset.

- ◆ **Recovery Interval** – Specifies the period after which to automatically recover from UDLD disabled port state. (Range: 30-86400 seconds; Default: 7 seconds)

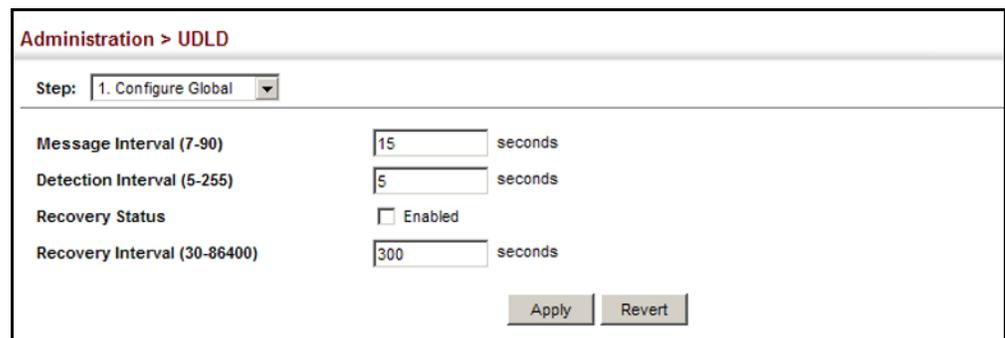
When the recovery interval is changed, any ports shut down by UDLD will be reset.

Web Interface

To configure the UDLD message probe interval, detection interval, and recovery interval:

1. Click Administration, UDLD, Configure Global.
2. Select Configure Global from the Step list.
3. Configure the message and detection intervals.
4. Enable automatic recovery if required, and set the recovery interval.
5. Click Apply.

Figure 267: Configuring UDLD Protocol Intervals



The screenshot shows the 'Administration > UDLD' configuration page. At the top, there is a breadcrumb trail 'Administration > UDLD' and a 'Step:' dropdown menu set to '1. Configure Global'. Below this, there are four configuration rows:

Message Interval (7-90)	<input type="text" value="15"/>	seconds
Detection Interval (5-255)	<input type="text" value="5"/>	seconds
Recovery Status	<input type="checkbox"/> Enabled	
Recovery Interval (30-86400)	<input type="text" value="300"/>	seconds

At the bottom right of the form, there are two buttons: 'Apply' and 'Revert'.

Configuring UDLD Interface Settings Use the Administration > UDLD (Configure Interface) page to enable UDLD and aggressive mode which reduces the shut-down delay after loss of bidirectional connectivity is detected.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/54)
- ◆ **UDLD** – Enables UDLD on a port. (Default: Disabled)
 - UDLD requires that all the devices connected to the same LAN segment be running the protocol in order for a potential mis-configuration to be detected and for prompt corrective action to be taken.
 - Whenever a UDLD device learns about a new neighbor or receives a resynchronization request from an out-of-synch neighbor, it (re)starts the detection process on its side of the connection and sends N echo messages in reply. (This mechanism implicitly assumes that N packets are sufficient to get through a link and reach the other end, even though some of them might get dropped during the transmission.)

Since this behavior must be the same on all the neighbors, the sender of the echoes expects to receive an echo in reply. If the detection process ends without the proper echo information being received, the link is considered to be unidirectional.
- ◆ **Aggressive Mode** – Reduces the shut-down delay after loss of bidirectional connectivity is detected. (Default: Disabled)

UDLD can function in two modes: normal mode and aggressive mode.

- In normal mode, determination of link status at the end of the detection process is always based on information received in UDLD messages: whether that's information about the exchange of proper neighbor identification or the absence of such. Hence, albeit bound by a timer, normal mode determinations are always based on gleaned information, and as such are "event-based." If no such information can be obtained (e.g., because of a bidirectional loss of connectivity), UDLD follows a conservative approach to minimize false positives during the detection process and deems a port to be in "undetermined" state. In other words, normal mode will shut down a port only if it can explicitly determine that the associated link is faulty for an extended period of time.
- In aggressive mode, UDLD will also shut down a port if it loses bidirectional connectivity with the neighbor for the same extended period of time (as that mentioned above for normal mode) and subsequently fails repeated last-resort attempts to re-establish communication with the other end of the link. This mode of operation assumes that loss of communication with the neighbor is a meaningful network event in itself, and a symptom of a serious connectivity problem. Because this type of detection can be event-less, and lack of information cannot always be associated to an actual malfunction of the link, this mode is recommended only in certain

scenarios (typically only on point-to-point links where no communication failure between two neighbors is admissible).

- ◆ **Operation State** – Shows the UDLD operational state (Disabled, Link down, Link up, Advertisement, Detection, Disabled port, Advertisement - Single neighbor, Advertisement - Multiple neighbors)

- ◆ **Port State** – Shows the UDLD port state (Unknown, Bidirectional, Unidirectional, Transmit-to-receive loop, Mismatch with neighbor state reported, Neighbor's echo is empty)

The state is Unknown if the link is down or not connected to a UDLD-capable device. The state is Bidirectional if the link has a normal two-way connection to a UDLD-capable device. All other states indicate mis-wiring.

- ◆ **Message Interval** – The interval between UDLD probe messages used for the indicated operational state.

- ◆ **Detection Interval** – The period the switch remains in detection state after discovering a neighbor.

Web Interface

To enable UDLD and aggressive mode:

1. Click Administration, UDLD, Configure Interface.
2. Enable UDLD and aggressive mode on the required ports.
3. Click Apply.

Figure 268: Configuring UDLD Interface Settings

Port	UDLD	Aggressive Mode	Operation State	Port State	Message Interval (seconds)	Detection Interval (seconds)
1	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	Disabled	Unknown	7	5
2	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	Disabled	Unknown	7	5
3	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	Disabled	Unknown	7	5
4	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	Disabled	Unknown	7	5
5	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	Disabled	Unknown	7	5

Displaying UDLD Neighbor Information

Use the Administration > UDLD (Show Information) page to show UDLD neighbor information, including neighbor state, expiration time, and protocol intervals.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1–32/54)

- ◆ **Entry** – Table entry number uniquely identifying the neighbor device discovered by UDLD on a port interface.
- ◆ **Device ID** – Device identifier of neighbor sending the UDLD packet.
- ◆ **Port ID** – The physical port the UDLD packet is sent from.
- ◆ **Device Name** – The device name of this neighbor.
- ◆ **Neighbor State** – Link status of neighbor device (Values: unknown, neighborsEcholsEmpty, bidirectional, mismatchWithneighborStateReported, unidirectional).
- ◆ **Expire** – The amount of time remaining before this entry will expire.
- ◆ **Message Interval** – The interval between UDLD probe messages for ports in advertisement phase.
- ◆ **Detection Interval** – The period the switch remains in detection state after discovering a neighbor.

Web Interface

To display UDLD neighbor information:

1. Click Administration, UDLD, Show Information.
2. Select an interface from the Port list.

Figure 269: Displaying UDLD Neighbor Information

Entry	Device ID	Port ID	Device Name	Neighbor State	Expire (seconds)	Message Interval (seconds)	Detection Interval (seconds)
1	S123456	Eth 1/1	ECS4110-S2P	Bidirectional	21	7	5

13

Multicast Filtering

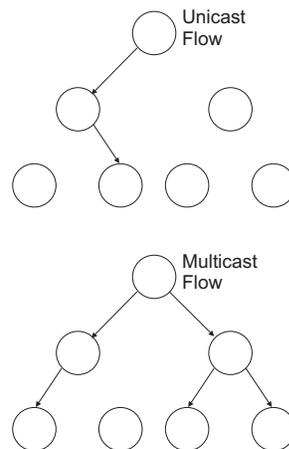
This chapter describes how to configure the following multicast services:

- ◆ **IGMP Snooping** – Configures snooping and query parameters for IPv4.
- ◆ **Filtering and Throttling** – Filters specified multicast service, or throttles the maximum of multicast groups allowed on an interface for IPv4.
- ◆ **MLD Snooping** – Configures snooping and query parameters for IPv6.
- ◆ **Layer 3 IGMP** – Configures IGMP query used with multicast routing.

Overview

Multicasting is used to support real-time applications such as video conferencing or streaming audio. A multicast server does not have to establish a separate connection with each client. It merely broadcasts its service to the network, and any hosts that want to receive the multicast register with their local multicast switch/router. Although this approach reduces the network overhead required by a multicast server, the broadcast traffic must be carefully pruned at every multicast switch/router it passes through to ensure that traffic is only passed on to the hosts which subscribed to this service.

Figure 270: Multicast Filtering Concept



This switch can use Internet Group Management Protocol (IGMP) to filter multicast traffic. IGMP Snooping can be used to passively monitor or “snoop” on exchanges between attached hosts and an IGMP-enabled device, most commonly a multicast

router. In this way, the switch can discover the ports that want to join a multicast group, and set its filters accordingly.

If there is no multicast router attached to the local subnet, multicast traffic and query messages may not be received by the switch. In this case (Layer 2) IGMP Query can be used to actively ask the attached hosts if they want to receive a specific multicast service. IGMP Query thereby identifies the ports containing hosts requesting to join the service and sends data out to those ports only. It then propagates the service request up to any neighboring multicast switch/router to ensure that it will continue to receive the multicast service.

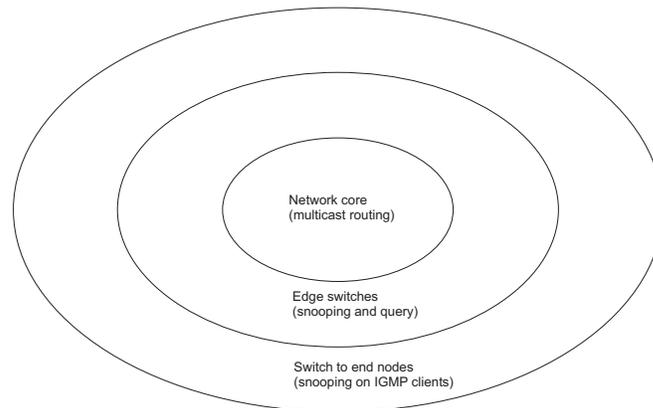
The purpose of IP multicast filtering is to optimize a switched network's performance, so multicast packets will only be forwarded to those ports containing multicast group hosts or multicast routers/switches, instead of flooding traffic to all ports in the subnet (VLAN).

This switch not only supports IP multicast filtering by passively monitoring IGMP query, report messages and multicast routing probe messages to register end-stations as multicast group members, but also supports the Protocol Independent Multicasting (PIM) routing protocol required to forward multicast traffic to other subnets ([page 607](#) and [623](#)).

IGMP Protocol

The Internet Group Management Protocol (IGMP) runs between hosts and their immediately adjacent multicast router/switch. IGMP is a multicast host registration protocol that allows any host to inform its local router that it wants to receive transmissions addressed to a specific multicast group. A router, or multicast-enabled switch, can periodically ask their hosts if they want to receive multicast traffic. If there is more than one router/switch on the LAN performing IP multicasting, one of these devices is elected "querier" (at Layer 3) and assumes the role of querying the LAN for group members. It then propagates the service requests on to any adjacent multicast switch/router to ensure that it will continue to receive the multicast service. Based on the group membership information learned from IGMP, a router/switch can determine which (if any) multicast traffic needs to be forwarded to each of its ports. At Layer 3, multicast routers use this information, along with a multicast routing protocol such as Protocol Independent Multicasting (PIM), to support IP multicasting across the Internet. Note that IGMP neither alters nor routes IP multicast packets. A multicast routing protocol must be used to deliver IP multicast packets across different subnetworks. Therefore, when PIM routing is enabled for a subnet on the switch, IGMP is automatically enabled.

Figure 271: IGMP Protocol



Layer 2 IGMP (Snooping and Query for IPv4)

IGMP Snooping and Query – If multicast routing is not supported on other switches in your network, you can use IGMP Snooping and IGMP Query ([page 437](#)) to monitor IGMP service requests passing between multicast clients and servers, and dynamically configure the switch ports which need to forward multicast traffic. IGMP Snooping conserves bandwidth on network segments where no node has expressed interest in receiving a specific multicast service. For switches that do not support multicast routing, or where multicast routing is already enabled on other switches in the local network segment, IGMP Snooping is the only service required to support multicast filtering.

When using IGMPv3 snooping, service requests from IGMP Version 1, 2 or 3 hosts are all forwarded to the upstream router as IGMPv3 reports. The primary enhancement provided by IGMPv3 snooping is in keeping track of information about the specific multicast sources which downstream IGMPv3 hosts have requested or refused. The switch maintains information about both multicast groups and channels, where a group indicates a multicast flow for which the hosts have *not* requested a specific source (the only option for IGMPv1 and v2 hosts unless statically configured on the switch), and a channel indicates a flow for which the hosts have requested service from a specific source. For IGMPv1/v2 hosts, the source address of a channel is always null (indicating that any source is acceptable), but for IGMPv3 hosts, it may include a specific address when requested.

Only IGMPv3 hosts can request service from a specific multicast source. When downstream hosts request service from a specific source for a multicast service, these sources are all placed in the Include list, and traffic is forwarded to the hosts from each of these sources. IGMPv3 hosts may also request that service be forwarded from any source except for those specified. In this case, traffic is filtered from sources in the Exclude list, and forwarded from all other available sources.



Note: When the switch is configured to use IGMPv3 snooping, the snooping version may be downgraded to version 2 or version 1, depending on the version of the IGMP query packets detected on each VLAN.

Note: IGMP snooping will not function unless a multicast router port is enabled on the switch. This can be accomplished in one of two ways. A static router port can be manually configured (see [“Specifying Static Interfaces for an IPv4 Multicast Router” on page 441](#)). Using this method, the router port is never timed out, and will continue to function until explicitly removed. The other method relies on the switch to dynamically create multicast routing ports whenever multicast routing protocol packets or IGMP query packets are detected on a port.

Note: A maximum of up to 1024 multicast entries can be maintained for IGMP snooping [and 255 entries for Multicast Routing when both of these features are enabled](#). Once the table is full, no new entries are learned. Any subsequent multicast traffic not found in the table is dropped if unregistered-flooding is disabled (default behavior) and no router port is configured in the attached VLAN, or flooded throughout the VLAN if unregistered-flooding is enabled (see [“Configuring IGMP Snooping and Query Parameters” on page 437](#)).

Static IGMP Router Interface – If IGMP snooping cannot locate the IGMP querier, you can manually designate a known IGMP querier (i.e., a multicast router/switch) connected over the network to an interface on your switch ([page 441](#)). This interface will then join all the current multicast groups supported by the attached router/switch to ensure that multicast traffic is passed to all appropriate interfaces within the switch.

Static IGMP Host Interface – For multicast applications that you need to control more carefully, you can manually assign a multicast service to specific interfaces on the switch ([page 443](#)).

IGMP Snooping with Proxy Reporting – The switch supports last leave, and query suppression (as defined in DSL Forum TR-101, April 2006):

- ◆ When proxy reporting is disabled, all IGMP reports received by the switch are forwarded natively to the upstream multicast routers.
- ◆ Last Leave: Intercepts, absorbs and summarizes IGMP leaves coming from IGMP hosts. IGMP leaves are relayed upstream only when necessary, that is, when the last user leaves a multicast group.
- ◆ Query Suppression: Intercepts and processes IGMP queries in such a way that IGMP specific queries are never sent to client ports.

The only deviation from TR-101 is that ~~report suppression~~, and the marking of IGMP traffic initiated by the switch with priority bits as defined in R-250 is not supported.



Note: IGMP Query (Layer 2 or 3) – IGMP Query can be enabled globally at Layer 2, but can also be enabled for individual VLAN interfaces at Layer 3 (see [“Configuring](#)

IGMP Interface Parameters”). However, note that Layer 2 query is disabled if Layer 3 query is enabled.

Configuring IGMP Snooping and Query Parameters

Use the Multicast > IGMP Snooping > General page to configure the switch to forward multicast traffic intelligently. Based on the IGMP query and report messages, the switch forwards multicast traffic only to the ports that request it. This prevents the switch from broadcasting the traffic to all ports and possibly disrupting network performance.

Command Usage

◆ **IGMP Snooping** – This switch can passively snoop on IGMP Query and Report packets transferred between IP multicast routers/switches and IP multicast host groups to identify the IP multicast group members. It simply monitors the IGMP packets passing through it, picks out the group registration information, and configures the multicast filters accordingly.



Note: If unknown multicast traffic enters a VLAN which has been configured with a router port, the traffic is forwarded to that port. However, if no router port exists on the VLAN, the traffic is dropped if unregistered data flooding is disabled (default behavior), or flooded throughout the VLAN if unregistered data flooding is enabled (see “Unregistered Data Flooding” in the Command Attributes section).

◆ **IGMP Querier** – A router, or multicast-enabled switch, can periodically ask their hosts if they want to receive multicast traffic. If there is more than one router/switch on the LAN performing IP multicasting, one of these devices is elected “querier” and assumes the role of querying the LAN for group members. It then propagates the service requests on to any upstream multicast switch/router to ensure that it will continue to receive the multicast service.



Note: Multicast routers use this information from IGMP snooping and query reports, along with a multicast routing protocol such as PIM, to support IP multicasting across the Internet.

Parameters

These parameters are displayed:

◆ **IGMP Snooping Status** – When enabled, the switch will monitor network traffic to determine which hosts want to receive multicast traffic. This is referred to as IGMP Snooping. (Default: Disabled)

When IGMP snooping is enabled globally, the per VLAN interface settings for IGMP snooping take precedence (see “[Setting IGMP Snooping Status per Interface](#)” on page 445).

When IGMP snooping is disabled globally, snooping can still be configured per VLAN interface, but the interface settings will not take effect until snooping is re-enabled globally.

- ◆ **Proxy Reporting Status** – Enables IGMP Snooping with Proxy Reporting. (Default: Disabled)

When proxy reporting is enabled with this command, the switch performs “IGMP Snooping with Proxy Reporting” (as defined in DSL Forum TR-101, April 2006), including last leave, and query suppression.

Last leave sends out a proxy query when the last member leaves a multicast group, and query suppression means that specific queries are not forwarded from an upstream multicast router to hosts downstream from this device.

When proxy reporting is disabled, all IGMP reports received by the switch are forwarded natively to the upstream multicast routers.

- ◆ **TCN Flood** – Enables flooding of multicast traffic if a spanning tree topology change notification (TCN) occurs. (Default: Disabled)

When a spanning tree topology change occurs, the multicast membership information learned by switch may be out of date. For example, a host linked to one port before the topology change (TC) may be moved to another port after the change. To ensure that multicast data is delivered to all receivers, by default, a switch in a VLAN (with IGMP snooping enabled) that receives a Bridge Protocol Data Unit (BPDU) with TC bit set (by the root bridge) will enter into “multicast flooding mode” for a period of time until the topology has stabilized and the new locations of all multicast receivers are learned.

If a topology change notification (TCN) is received, and all the uplink ports are subsequently deleted, a time out mechanism is used to delete all of the currently learned multicast channels.

When a new uplink port starts up, the switch sends unsolicited reports for all currently learned channels out the new uplink port.

By default, the switch immediately enters into “multicast flooding mode” when a spanning tree topology change occurs. In this mode, multicast traffic will be flooded to all VLAN ports. If many ports have subscribed to different multicast groups, flooding may cause excessive packet loss on the link between the switch and the end host. Flooding may be disabled to avoid this, causing multicast traffic to be delivered only to those ports on which multicast group members have been learned. Otherwise, the time spent in flooding mode can be manually configured to reduce excessive loading.

When the spanning tree topology changes, the root bridge sends a proxy query to quickly re-learn the host membership/port relations for multicast channels. The root bridge also sends an unsolicited Multicast Router Discover (MRD) request to quickly locate the multicast routers in this VLAN.

The proxy query and unsolicited MRD request are flooded to all VLAN ports except for the receiving port when the switch receives such packets.

- ◆ **TCN Query Solicit** – Sends out an IGMP general query solicitation when a spanning tree topology change notification (TCN) occurs. (Default: Disabled)

When the root bridge in a spanning tree receives a TCN for a VLAN where IGMP snooping is enabled, it issues a global IGMP leave message (or query solicitation). When a switch receives this solicitation, it floods it to all ports in the VLAN where the spanning tree change occurred. When an upstream multicast router receives this solicitation, it immediately issues an IGMP general query.

A query solicitation can be sent whenever the switch notices a topology change, even if it is not the root bridge in spanning tree.
- ◆ **Router Alert Option** – Discards any IGMPv2/v3 packets that do not include the Router Alert option. (Default: Disabled)

As described in Section 9.1 of RFC 3376 for IGMP Version 3, the Router Alert Option can be used to protect against DOS attacks. One common method of attack is launched by an intruder who takes over the role of querier, and starts overloading multicast hosts by sending a large number of group-and-source-specific queries, each with a large source list and the Maximum Response Time set to a large value.

To protect against this kind of attack, (1) routers should not forward queries. This is easier to accomplish if the query carries the Router Alert option. (2) Also, when the switch is acting in the role of a multicast host (such as when using proxy routing), it should ignore version 2 or 3 queries that do not contain the Router Alert option.
- ◆ **Unregistered Data Flooding** – Floods unregistered multicast traffic into the attached VLAN. (Default: Disabled)

Once the table used to store multicast entries for IGMP snooping and multicast routing is filled, no new entries are learned. If no router port is configured in the attached VLAN, and unregistered-flooding is disabled, any subsequent multicast traffic not found in the table is dropped, otherwise it is flooded throughout the VLAN.
- ◆ **Forwarding Priority** – Assigns a CoS priority to all multicast traffic. (Range: 0-7, where 7 is the highest priority; Default: Disabled)

This parameter can be used to set a high priority for low-latency multicast traffic such as a video-conference, or to set a low priority for normal multicast traffic not sensitive to latency.
- ◆ **Version Exclusive** – Discards any received IGMP messages which use a version different to that currently configured by the IGMP Version attribute. (Default: Disabled)
- ◆ **IGMP Unsolicited Report Interval** – Specifies how often the upstream interface should transmit unsolicited IGMP reports when proxy reporting is enabled. (Range: 1-65535 seconds, Default: 400 seconds)

When a new upstream interface (that is, uplink port) starts up, the switch sends unsolicited reports for all currently learned multicast channels via the new upstream interface.

This command only applies when proxy reporting is enabled.

- ◆ **Router Port Expire Time** – The time the switch waits after the previous querier stops before it considers it to have expired. (Range: 1-65535 seconds; Default: 300 seconds)
- ◆ **IGMP Snooping Version** – Sets the protocol version for compatibility with other devices on the network. This is the IGMP Version the switch uses to send snooping reports. (Range: 1-3; Default: 2)

This attribute configures the IGMP report/query version used by IGMP snooping. Versions 1 - 3 are all supported, and versions 2 and 3 are backward compatible, so the switch can operate with other devices, regardless of the snooping version employed.

- ◆ **Querier Status** – When enabled, the switch can serve as the Querier, which is responsible for asking hosts if they want to receive multicast traffic. This feature is not supported for IGMPv3 snooping. (Default: Disabled)

Web Interface

To configure general settings for IGMP Snooping and Query:

1. Click Multicast, IGMP Snooping, General.
2. Adjust the IGMP settings as required.
3. Click Apply.

Figure 272: Configuring General Settings for IGMP Snooping

The screenshot shows a web interface titled "Multicast > IGMP Snooping > General". It contains a list of configuration options, each with a checkbox and a value field. The options are: IGMP Snooping Status (checkbox), Proxy Reporting Status (checkbox), TCN Flood (checkbox), TCN Query Solicit (checkbox), Router Alert Option (checkbox), Unregistered Data Flooding (checkbox), Forwarding Priority (0-7) (checkbox and input field), Version Exclusive (checkbox), IGMP Unsolicited Report Interval (1-65535) (input field with value 400 and label "seconds"), Router Port Expire Time (1-65535) (input field with value 300 and label "seconds"), IGMP Snooping Version (1-3) (input field with value 2), and Querier Status (checkbox). At the bottom right, there are "Apply" and "Revert" buttons.

Setting	Value
IGMP Snooping Status	<input type="checkbox"/> Enabled
Proxy Reporting Status	<input type="checkbox"/> Enabled
TCN Flood	<input type="checkbox"/> Enabled
TCN Query Solicit	<input type="checkbox"/> Enabled
Router Alert Option	<input type="checkbox"/> Enabled
Unregistered Data Flooding	<input type="checkbox"/> Enabled
Forwarding Priority (0-7)	<input type="checkbox"/> <input type="text"/>
Version Exclusive	<input type="checkbox"/> Enabled
IGMP Unsolicited Report Interval (1-65535)	<input type="text" value="400"/> seconds
Router Port Expire Time (1-65535)	<input type="text" value="300"/> seconds
IGMP Snooping Version (1-3)	<input type="text" value="2"/>
Querier Status	<input type="checkbox"/> Enabled

Specifying Static Interfaces for an IPv4 Multicast Router

Use the Multicast > IGMP Snooping > Multicast Router (Add Static Multicast Router) page to statically attach an IPv4 interface to a multicast router/switch.

Depending on network connections, IGMP snooping may not always be able to locate the IGMP querier. Therefore, if the IGMP querier is a known multicast router/switch connected over the network to an interface (port or trunk) on the switch, the interface (and a specified VLAN) can be manually configured to join all the current multicast groups supported by the attached router. This can ensure that multicast traffic is passed to all the appropriate interfaces within the switch.

Command Usage

IGMP Snooping must be enabled globally on the switch (see [“Configuring IGMP Snooping and Query Parameters” on page 437](#)) before a multicast router port can take effect.

Parameters

These parameters are displayed:

Add Static Multicast Router

- ◆ **VLAN** – Selects the VLAN which is to propagate all multicast traffic coming from the attached multicast router. (Range: 1-4094)
- ◆ **Interface** – Activates the Port or Trunk scroll down list.
- ◆ **Port or Trunk** – Specifies the interface attached to a multicast router.

Show Static Multicast Router

- ◆ **VLAN** – Selects the VLAN for which to display any configured static multicast routers.
- ◆ **Interface** – Shows the interface to which the specified static multicast routers are attached.

Show Current Multicast Router

- ◆ **VLAN** – Selects the VLAN for which to display any currently active multicast routers.
- ◆ **Interface** – Shows the interface to which an active multicast router is attached.
- ◆ **Type** – Shows if this entry is static or dynamic.
- ◆ **Expire** – Time until this dynamic entry expires.

Web Interface

To specify a static interface attached to a multicast router:

1. Click Multicast, IGMP Snooping, Multicast Router.
2. Select Add Static Multicast Router from the Action list.
3. Select the VLAN which will forward all the corresponding multicast traffic, and select the port or trunk attached to the multicast router.
4. Click Apply.

Figure 273: Configuring a Static Interface for an IPv4 Multicast Router

The screenshot shows the configuration page for a multicast router. The breadcrumb is "Multicast > IGMP Snooping > Multicast Router". The "Action" dropdown is set to "Add Static Multicast Router". The "VLAN" dropdown is set to "1". The "Interface" section has two radio buttons: "Port" (selected) and "Trunk". The "Port" radio button is followed by a dropdown menu set to "1". There are "Apply" and "Revert" buttons at the bottom right.

To show the static interfaces attached to a multicast router:

1. Click Multicast, IGMP Snooping, Multicast Router.
2. Select Show Static Multicast Router from the Action list.
3. Select the VLAN for which to display this information.

Figure 274: Showing Static Interfaces Attached an IPv4 Multicast Router

The screenshot shows the configuration page for a multicast router. The breadcrumb is "Multicast > IGMP Snooping > Multicast Router". The "Action" dropdown is set to "Show Static Multicast Router". The "VLAN" dropdown is set to "1". Below the dropdowns is a table titled "Static Multicast Router Interface List" with a "Total: 6" count. The table has a checkbox column and an "Interface" column. The interfaces listed are: Unit 1 / Port 1, Unit 1 / Port 2, Unit 1 / Port 3, Trunk 2, Trunk 5, and Unit 1 / Port 4. There are "Delete" and "Revert" buttons at the bottom right.

<input type="checkbox"/>	Interface
<input type="checkbox"/>	Unit 1 / Port 1
<input type="checkbox"/>	Unit 1 / Port 2
<input type="checkbox"/>	Unit 1 / Port 3
<input type="checkbox"/>	Trunk 2
<input type="checkbox"/>	Trunk 5
<input type="checkbox"/>	Unit 1 / Port 4

Multicast routers that are attached to ports on the switch use information obtained from IGMP, along with a multicast routing protocol (such as PIM) to support IP multicasting across the Internet. These routers may be dynamically discovered by

the switch or statically assigned to an interface on the switch. To show all the interfaces attached to a multicast router:

1. Click Multicast, IGMP Snooping, Multicast Router.
2. Select Show Current Multicast Router from the Action list.
3. Select the VLAN for which to display this information. Ports in the selected VLAN which are attached to a neighboring multicast router/switch are displayed.

Figure 275: Showing Current Interfaces Attached an IPv4 Multicast Router

The screenshot shows a web interface for configuring Multicast Router. The breadcrumb path is 'Multicast > IGMP Snooping > Multicast Router'. The 'Action' dropdown is set to 'Show Current Multicast Router'. The 'VLAN' dropdown is set to '1'. Below this, there is a table titled 'Multicast Router Interface Information' with a total of 2 entries.

Interface	Type	Expire
Eth 1 / 1	Static	
Eth 1 / 4	Dynamic	0:4:59

Assigning Interfaces to IPv4 Multicast Services

Use the Multicast > IGMP Snooping > IGMP Member (Add Static Member) page to statically assign an IPv4 multicast service to an interface.

Multicast filtering can be dynamically configured using IGMP Snooping and IGMP Query messages (see [“Configuring IGMP Snooping and Query Parameters” on page 437](#)). However, for certain applications that require tighter control, it may be necessary to statically configure a multicast service on the switch. First add all the ports attached to participating hosts to a common VLAN, and then assign the multicast service to that VLAN group.

Command Usage

- ◆ Static multicast addresses are never aged out.
- ◆ When a multicast address is assigned to an interface in a specific VLAN, the corresponding traffic can only be forwarded to ports within that VLAN.

Parameters

These parameters are displayed:

- ◆ **VLAN** – Specifies the VLAN which is to propagate the multicast service. (Range: 1-4094)
- ◆ **Interface** – Activates the Port or Trunk scroll down list.
- ◆ **Port or Trunk** – Specifies the interface assigned to a multicast group.
- ◆ **Multicast IP** – The IP address for a specific multicast service.

Web Interface

To statically assign an interface to an IPv4 multicast service:

1. Click Multicast, IGMP Snooping, IGMP Member.
2. Select Add Static Member from the Action list.
3. Select the VLAN that will propagate the multicast service, specify the interface attached to a multicast service (through an IGMP-enabled switch or multicast router), and enter the multicast IP address.
4. Click Apply.

Figure 276: Assigning an Interface to an IPv4 Multicast Service

The screenshot shows the 'Multicast > IGMP Snooping > IGMP Member' configuration page. The 'Action' dropdown is set to 'Add Static Member'. The 'VLAN' dropdown is set to '1'. The 'Interface' section has 'Port 1' selected with a radio button, and 'Trunk 1' is also visible. The 'Multicast IP' text field contains '224.1.1.1'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show the static interfaces assigned to an IPv4 multicast service:

1. Click Multicast, IGMP Snooping, IGMP Member.
2. Select Show Static Member from the Action list.
3. Select the VLAN for which to display this information.

Figure 277: Showing Static Interfaces Assigned to an IPv4 Multicast Service

The screenshot shows the 'Multicast > IGMP Snooping > IGMP Member' configuration page with the 'Action' dropdown set to 'Show Static Member'. The 'VLAN' dropdown is set to '1'. Below the dropdowns, it says 'IGMP Member Interface List Total: 6'. A table displays the following data:

<input type="checkbox"/>	Interface	Multicast IP
<input type="checkbox"/>	Unit 1 / Port 1	224.1.1.1
<input type="checkbox"/>	Unit 1 / Port 2	224.1.2.2
<input type="checkbox"/>	Unit 1 / Port 3	230.1.1.1
<input type="checkbox"/>	Trunk 2	230.1.2.2
<input type="checkbox"/>	Trunk 5	239.1.1.1
<input type="checkbox"/>	Unit 1 / Port 4	239.2.2.2

At the bottom right, there are 'Delete' and 'Revert' buttons.

To show the all interfaces attached to a multicast router:

1. Click Multicast, IGMP Snooping, Multicast Router.
2. Select Current Multicast Router from the Action list.
3. Select the VLAN for which to display this information. Ports in the selected VLAN which are attached to a neighboring multicast router/switch are displayed.

Figure 278: Showing Current Interfaces Attached a Multicast Router

The screenshot shows a web interface for configuring IGMP Snooping. The breadcrumb is 'Multicast > IGMP Snooping > Multicast Router'. The 'Action' dropdown is set to 'Show Current Multicast Router'. The 'VLAN' dropdown is set to '1'. Below this, a table titled 'Multicast Router Interface Information' shows a total of 2 interfaces. The table has three columns: 'Interface', 'Type', and 'Expire'.

Interface	Type	Expire
Eth 1 / 1	Static	
Eth 1 / 4	Dynamic	0:4:59

Setting IGMP Snooping Status per Interface

Use the Multicast > IGMP Snooping > Interface (Configure VLAN) page to configure IGMP snooping attributes for a VLAN. To configure snooping globally, refer to [“Configuring IGMP Snooping and Query Parameters”](#) on page 437.

Command Usage

Multicast Router Discovery

There have been many mechanisms used in the past to identify multicast routers. This has lead to interoperability issues between multicast routers and snooping switches from different vendors. In response to this problem, the Multicast Router Discovery (MRD) protocol has been developed for use by IGMP snooping and multicast routing devices. MRD is used to discover which interfaces are attached to multicast routers, allowing IGMP-enabled devices to determine where to send multicast source and group membership messages. (MRD is specified in draft-ietf-magma-mrdisc-07.)

Multicast source data and group membership reports must be received by all multicast routers on a segment. Using the group membership protocol query messages to discover multicast routers is insufficient due to query suppression. MRD therefore provides a standardized way to identify multicast routers without relying on any particular multicast routing protocol.



Note: The default values recommended in the MRD draft are implemented in the switch.

Multicast Router Discovery uses the following three message types to discover multicast routers:

- ◆ **Multicast Router Advertisement** – Advertisements are sent by routers to advertise that IP multicast forwarding is enabled. These messages are sent unsolicited periodically on all router interfaces on which multicast forwarding is enabled. They are sent upon the occurrence of these events:
 - Upon the expiration of a periodic (randomized) timer.
 - As a part of a router's start up procedure.
 - During the restart of a multicast forwarding interface.
 - On receipt of a Solicitation message.
- ◆ **Multicast Router Solicitation** – Devices send Solicitation messages in order to solicit Advertisement messages from multicast routers. These messages are used to discover multicast routers on a directly attached link. Solicitation messages are also sent whenever a multicast forwarding interface is initialized or re-initialized. Upon receiving a solicitation on an interface with IP multicast forwarding and MRD enabled, a router will respond with an Advertisement.
- ◆ **Multicast Router Termination** – These messages are sent when a router stops IP multicast routing functions on an interface. Termination messages are sent by multicast routers when:
 - Multicast forwarding is disabled on an interface.
 - An interface is administratively disabled.
 - The router is gracefully shut down.

Advertisement and Termination messages are sent to the All-Snoopers multicast address. Solicitation messages are sent to the All-Routers multicast address.



Note: MRD messages are flooded to all ports in a VLAN where IGMP snooping or routing has been enabled. To ensure that older switches which do not support MRD can also learn the multicast router port, the switch floods IGMP general query packets, which do not have a null source address (0.0.0.0), to all ports in the attached VLAN. IGMP packets with a null source address are only flooded to all ports in the VLAN if the system is operating in multicast flooding mode, such as when a new VLAN or new router port is being established, or a spanning tree topology change has occurred. Otherwise, this kind of packet is only forwarded to known multicast routing ports.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of configured VLANs. (Range: 1-4094)
- ◆ **IGMP Snooping Status** – When enabled, the switch will monitor network traffic on the indicated VLAN interface to determine which hosts want to

receive multicast traffic. This is referred to as IGMP Snooping.
(Default: Disabled)

When IGMP snooping is enabled globally (see [page 437](#)), the per VLAN interface settings for IGMP snooping take precedence.

When IGMP snooping is disabled globally, snooping can still be configured per VLAN interface, but the interface settings will not take effect until snooping is re-enabled globally.

- ◆ **Version Exclusive** – Discards any received IGMP messages (except for multicast protocol packets) which use a version different to that currently configured by the IGMP Version attribute. (Options: Enabled, Using Global Status; Default: Using Global Status)

If version exclusive is disabled on a VLAN, then this setting is based on the global setting configured on the Multicast > IGMP Snooping > General page. If it is enabled on a VLAN, then this setting takes precedence over the global setting.

- ◆ **Immediate Leave Status** – Immediately deletes a member port of a multicast service if a leave packet is received at that port and immediate leave is enabled for the parent VLAN. (Default: Disabled)

If immediate leave is not used, a multicast router (or querier) will send a group-specific query message when an IGMPv2 group leave message is received. The router/querier stops forwarding traffic for that group only if no host replies to the query within the specified time out period. Note that this time out is set to Last Member Query Interval * Robustness Variable (fixed at 2 as defined in RFC 2236).

If immediate leave is enabled, the switch assumes that only one host is connected to the interface. Therefore, immediate leave should only be enabled on an interface if it is connected to only one IGMP-enabled device, either a service host or a neighbor running IGMP snooping.

This attribute is only effective if IGMP snooping is enabled, and IGMPv2 snooping is used.

If immediate leave is enabled, the following options are provided:

- **By Group** – The switch assumes that only one host is connected to the interface. Therefore, immediate leave should only be enabled on an interface if it is connected to only one IGMP-enabled device, either a service host or a neighbor running IGMP snooping.
 - **By Host IP** – The switch will not send out a group-specific query when an IGMPv2/v3 leave message is received. But will check if there are other hosts joining the multicast group. Only when all hosts on that port leave the group will the member port be deleted.
- ◆ **Multicast Router Discovery** – MRD is used to discover which interfaces are attached to multicast routers. (Default: Disabled)

- ◆ **General Query Suppression** – Suppresses general queries except for ports attached to downstream multicast hosts. (Default: Disabled)

By default, general query messages are flooded to all ports, except for the multicast router through which they are received.

If general query suppression is enabled, then these messages are forwarded only to downstream ports which have joined a multicast service.

- ◆ **Proxy Reporting** – Enables IGMP Snooping with Proxy Reporting. (Options: Enabled, Disabled, Using Global Status; Default: Using Global Status)

When proxy reporting is enabled with this command, the switch performs “IGMP Snooping with Proxy Reporting” (as defined in DSL Forum TR-101, April 2006), including last leave, and query suppression.

Last leave sends out a proxy query when the last member leaves a multicast group, and query suppression means that specific queries are not forwarded from an upstream multicast router to hosts downstream from this device.

Rules Used for Proxy Reporting

When IGMP Proxy Reporting is disabled, the switch will use a null IP address for the source of IGMP query and report messages unless a proxy query address has been set.

When IGMP Proxy Reporting is enabled, the source address is based on the following criteria:

- If a proxy query address is configured, the switch will use that address as the source IP address in general and group-specific query messages sent to downstream hosts, and in report and leave messages sent upstream from the multicast router port.
 - If a proxy query address is not configured, the switch will use the VLAN's IP address as the IP source address in general and group-specific query messages sent downstream, and use the source address of the last IGMP message received from a downstream host in report and leave messages sent upstream from the multicast router port.
- ◆ **Interface Version** – Sets the protocol version for compatibility with other devices on the network. This is the IGMP Version the switch uses to send snooping reports. (Options: 1-3, Using Global Version; Default: Using Global Version)
- ◆ **Query Interval** – The interval between sending IGMP proxy general queries. (Range: 2-31744 seconds; Default: 125 seconds)

This attribute configures the IGMP report/query version used by IGMP snooping. Versions 1 - 3 are all supported, and versions 2 and 3 are backward compatible, so the switch can operate with other devices, regardless of the snooping version employed.

An IGMP general query message is sent by the switch at the interval specified by this attribute. When this message is received by downstream hosts, all receivers build an IGMP report for the multicast groups they have joined.

This attribute applies when the switch is serving as the querier (page 437), or as a proxy host when IGMP snooping proxy reporting is enabled (page 437).

- ◆ **Query Response Interval** – The maximum time the system waits for a response to proxy general queries. (Range: 10-31740 tenths of a second in multiples of 10; Default: 10 seconds)

This attribute applies when the switch is serving as the querier (page 437), or as a proxy host when IGMP snooping proxy reporting is enabled (page 437).

- ◆ **Last Member Query Interval** – The interval to wait for a response to a group-specific or group-and-source-specific query message. (Range: 1-31744 tenths of a second in multiples of 10; Default: 1 second)

When a multicast host leaves a group, it sends an IGMP leave message. When the leave message is received by the switch, it checks to see if this host is the last to leave the group by sending out an IGMP group-specific or group-and-source-specific query message, and starts a timer. If no reports are received before the timer expires, the group record is deleted, and a report is sent to the upstream multicast router.

A reduced value will result in reduced time to detect the loss of the last member of a group or source, but may generate more burst traffic.

This attribute will take effect only if IGMP snooping proxy reporting is enabled (see page 437) or IGMP querier is enabled (page 437).

- ◆ **Last Member Query Count** – The number of IGMP proxy group-specific or group-and-source-specific query messages that are sent out before the system assumes there are no more local members. (Range: 1-255; Default: 2)

This attribute will take effect only if IGMP snooping proxy reporting or IGMP querier is enabled.

- ◆ **Proxy Query Address** – A static source address for locally generated query and report messages used by IGMP Proxy Reporting. (Range: Any valid IP unicast address; Default: 0.0.0.0)

IGMP Snooping uses a null IP address of 0.0.0.0 for the source of IGMP query messages which are proxied to downstream hosts to indicate that it is not the elected querier, but is only proxying these messages as defined in RFC 4541. The switch also uses a null address in IGMP reports sent to upstream ports.

Many hosts do not implement RFC 4541, and therefore do not understand query messages with the source address of 0.0.0.0. These hosts will therefore not reply to the queries, causing the multicast router to stop sending traffic to them.

To resolve this problem, the source address in proxied IGMP query messages can be replaced with any valid unicast address (other than the router's own address).

Web Interface

To configure IGMP snooping on a VLAN:

1. Click Multicast, IGMP Snooping, Interface.
2. Select Configure VLAN from the Action list.
3. Select the VLAN to configure and update the required parameters.
4. Click Apply.

Figure 279: Configuring IGMP Snooping on a VLAN

Multicast > IGMP Snooping > Interface

Action: Configure VLAN

VLAN: 1

IGMP Snooping Status: Enabled

Version Exclusive: Using Global Status

Immediate Leave Status: Enabled By-Group

Multicast Router Discovery: Enabled

General Query Suppression: Enabled

Proxy Reporting: Using Global Status

Interface Version: Using Global Version

Query Interval (2-31744): seconds

Query Response Interval (10-31740): (1/10 seconds, multiple of 10)

Last Member Query Interval (1-31744): (1/10 seconds, multiple of 10)

Last Member Query Count (1-255):

Proxy (Query) Address:

Apply Revert

To show the interface settings for IGMP snooping:

1. Click Multicast, IGMP Snooping, Interface.
2. Select Show VLAN Information from the Action list.

Figure 280: Showing Interface Settings for IGMP Snooping

The screenshot shows the configuration page for IGMP Snooping on an interface. The breadcrumb is "Multicast > IGMP Snooping > Interface". Below the breadcrumb, there is an "Action:" dropdown menu set to "Show VLAN Information". The main content is a table titled "IGMP Snooping VLAN List" with a "Total: 4" indicator. The table has 13 columns: VLAN, IGMP Snooping Status, Immediate Leave Status, Query Interval, Query Response Interval, Last Member Query Interval, Last Member Query Count, Proxy (Query) Address, Proxy Reporting, Multicast Router Discovery, General Query Suppression, Version Exclusive, and Interface Version. The table contains four rows of data for VLANs 1, 2, 3, and 10.

VLAN	IGMP Snooping Status	Immediate Leave Status	Query Interval	Query Response Interval	Last Member Query Interval	Last Member Query Count	Proxy (Query) Address	Proxy Reporting	Multicast Router Discovery	General Query Suppression	Version Exclusive	Interface Version
1	Enabled	Disabled	10	100	10	2	10.1.1.1	Enabled	Enabled	Disabled	Enabled	1
2	Disabled	Disabled	10	100	10	2	20.2.2.2	Disabled	Disabled	Enabled	Disabled	3
3	Disabled	Disabled	10	100	10	2	30.3.3.3	Disabled	Enabled	Disabled	Disabled	2
10	Disabled	Disabled	10	100	10	2	100.10.10.10	Disabled	Disabled	Enabled	Disabled	1

Filtering IGMP Query Packets Use the Multicast > IGMP Snooping > Interface (Configure Interface) page to configure an interface to drop IGMP query packets.

Parameters

These parameters are displayed:

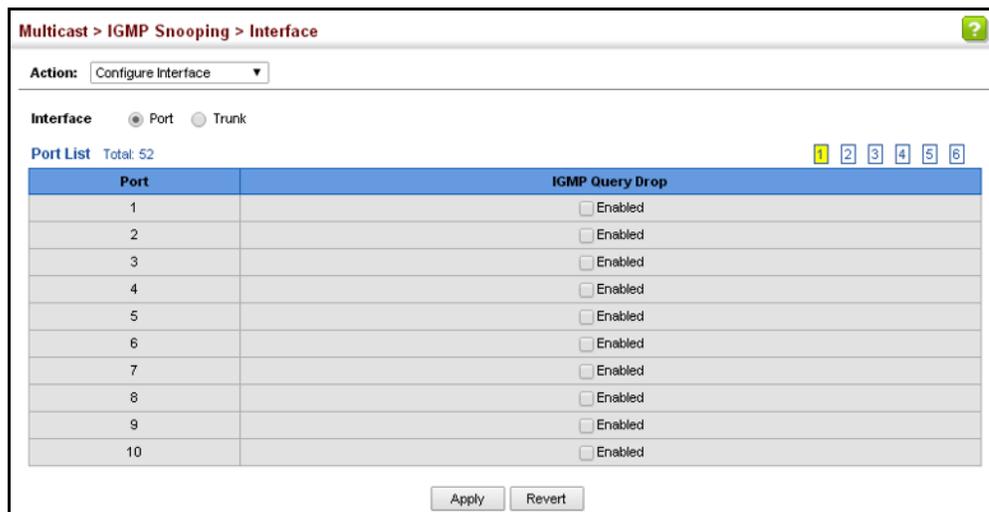
- ◆ **Interface** – Port or trunk identifier.
- ◆ **IGMP Query Drop** – Configures an interface to drop any IGMP query packets received on the specified interface. If this switch is acting as a Querier, this prevents it from being affected by messages received from another Querier.

Web Interface

To drop IGMP query packets or multicast data packets:

1. Click Multicast, IGMP Snooping, Interface.
2. Select Configure Interface from the Action List.
3. Click Port or Trunk to display the required interface type.
4. Enable the required drop functions for any interface.
5. Click Apply.

Figure 281: Dropping IGMP Query Packets



Displaying Multicast Groups Discovered by IGMP Snooping

Use the Multicast > IGMP Snooping > Forwarding Entry page to display the forwarding entries learned through IGMP Snooping.

Command Usage

To display information about multicast groups, IGMP Snooping must first be enabled on the switch (see [page 437](#)).

Parameters

These parameters are displayed:

- ◆ **VLAN** – An interface on the switch that is forwarding traffic to downstream ports for the specified multicast group address.
- ◆ **Group Address** – IP multicast group address with subscribers directly attached or downstream from the switch, or a static multicast group assigned to this interface.
- ◆ **Source Address** – The address of one of the multicast servers transmitting traffic to the specified group.
- ◆ **Interface** – A downstream port or trunk that is receiving traffic for the specified multicast group. This field may include both dynamically and statically configured multicast router ports.
- ◆ **Up Time** – Time that this multicast group has been known.
- ◆ **Expire** – The time until this entry expires.
- ◆ **Count** – The number of times this address has been learned by IGMP snooping.

Web Interface

To show multicast groups learned through IGMP snooping:

1. Click Multicast, IGMP Snooping, Forwarding Entry.
2. Select the VLAN for which to display this information.

Figure 282: Showing Multicast Groups Learned by IGMP Snooping

Group Address	Source Address	Interface	Up Time	Expire	Count
224.1.1.1	*	Eth 1 / 9 (Router Port)	00:00:06:46		2 (Port)
		Eth 1 / 11 (Member Port)	00:00:06:46	03:46	1 (Host)
224.1.1.2	192.168.1.2	Eth 1 / 9 (Router Port)		02:24	1 (Port)
224.1.1.3	*	Eth 1 / 9 (Router Port)	00:00:16:14		1 (Port)
239.255.255.250	*	Eth 1 / 9 (Router Port)	00:00:08:47		2 (Port)
		Eth 1 / 11 (Member Port)	00:00:08:47	03:46	1 (Host)

Displaying IGMP Snooping Statistics Use the Multicast > IGMP Snooping > Statistics pages to display IGMP snooping protocol-related statistics for the specified interface.

Parameters

These parameters are displayed:

- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Port** – Port identifier. (Range: 1-32/54)
- ◆ **Trunk** – Trunk identifier. (Range: 1-27)

Query Statistics

- ◆ **Other Querier** – IP address of remote querier on this interface.
- ◆ **Other Querier Expire** – Time after which remote querier is assumed to have expired.
- ◆ **Other Querier Uptime** – Time remote querier has been up.
- ◆ **Self Querier** – IP address of local querier on this interface.
- ◆ **Self Querier Expire** – Time after which local querier is assumed to have expired.
- ◆ **Self Querier Uptime** – Time local querier has been up.

- ◆ **General Query Received** – The number of general queries received on this interface.
- ◆ **General Query Sent** – The number of general queries sent from this interface.
- ◆ **Specific Query Received** – The number of specific queries received on this interface.
- ◆ **Specific Query Sent** – The number of specific queries sent from this interface.
- ◆ **Warn Rate Limit** – The rate at which received query messages of the wrong version type cause the Vx warning count to increment. Note that “0 sec” means that the Vx warning count is incremented for each wrong message version received.
- ◆ **V1 Warning Count** – The number of times the query version received (Version 1) does not match the version configured for this interface.
- ◆ **V2 Warning Count** – The number of times the query version received (Version 2) does not match the version configured for this interface.
- ◆ **V3 Warning Count** – The number of times the query version received (Version 3) does not match the version configured for this interface.

VLAN, Port, and Trunk Statistics

Input Statistics

- ◆ **Report** – The number of IGMP membership reports received on this interface.
- ◆ **Leave** – The number of leave messages received on this interface.
- ◆ **G Query** – The number of general query messages received on this interface.
- ◆ **G(-S)-S Query** – The number of group specific or group-and-source specific query messages received on this interface.
- ◆ **Drop** – The number of times a report, leave or query was dropped. Packets may be dropped due to invalid format, rate limiting, packet content not allowed, or IGMP group report received.
- ◆ **Join Success** – The number of times a multicast group was successfully joined.
- ◆ **Group** – The number of IGMP groups active on this interface.

Output Statistics

- ◆ **Report** – The number of IGMP membership reports sent from this interface.
- ◆ **Leave** – The number of leave messages sent from this interface.

- ◆ **G Query** – The number of general query messages sent from this interface.
- ◆ **G(-S)-S Query** – The number of group specific or group-and-source specific query messages sent from this interface.

Web Interface

To display statistics for IGMP snooping query-related messages:

1. Click Multicast, IGMP Snooping, Statistics.
2. Select Show Query Statistics from the Action list.
3. Select a VLAN.

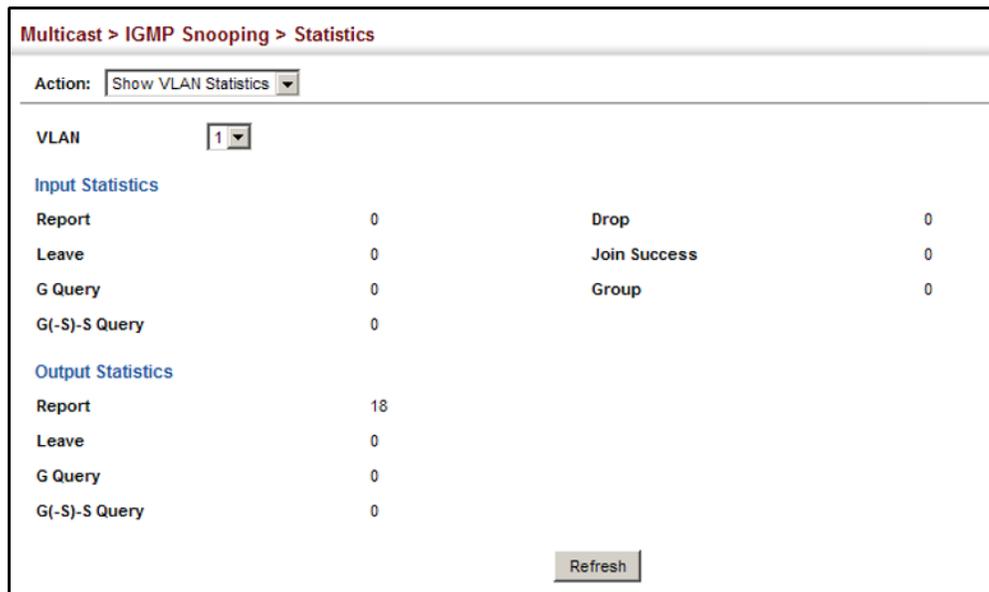
Figure 283: Displaying IGMP Snooping Statistics – Query

Query Statistics	
Other Querier	None
Other Querier Expire	00(m):00(s)
Other Querier Uptime	00(h):00(m):00(s)
Self Querier	192.168.1.1
Self Querier Expire	00(m):00(s)
Self Querier Uptime	00(h):00(m):00(s)
General Query Received	0
General Query Sent	0
Specific Query Received	0
Specific Query Sent	0
Warn Rate Limit	0 sec.
V1 Warning Count	0
V2 Warning Count	0
V3 Warning Count	0

To display IGMP snooping protocol-related statistics for a VLAN:

1. Click Multicast, IGMP Snooping, Statistics.
2. Select Show VLAN Statistics from the Action list.
3. Select a VLAN.

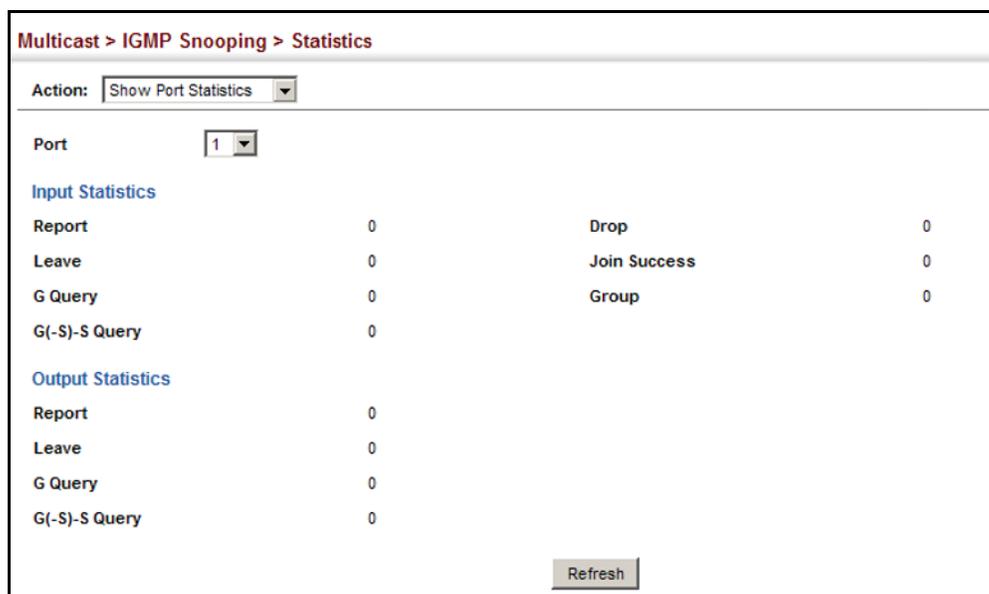
Figure 284: Displaying IGMP Snooping Statistics – VLAN



To display IGMP snooping protocol-related statistics for a port:

1. Click Multicast, IGMP Snooping, Statistics.
2. Select Show Port Statistics from the Action list.
3. Select a Port.

Figure 285: Displaying IGMP Snooping Statistics – Port



Filtering and Throttling IGMP Groups

In certain switch applications, the administrator may want to control the multicast services that are available to end users. For example, an IP/TV service based on a specific subscription plan. The IGMP filtering feature fulfills this requirement by restricting access to specified multicast services on a switch port, and IGMP throttling limits the number of simultaneous multicast groups a port can join.

IGMP filtering enables you to assign a profile to a switch port that specifies multicast groups that are permitted or denied on the port. An IGMP filter profile can contain one or more addresses, or a range of multicast addresses; but only one profile can be assigned to a port. When enabled, IGMP join reports received on the port are checked against the filter profile. If a requested multicast group is permitted, the IGMP join report is forwarded as normal. If a requested multicast group is denied, the IGMP join report is dropped.

IGMP throttling sets a maximum number of multicast groups that a port can join at the same time. When the maximum number of groups is reached on a port, the switch can take one of two actions; either “deny” or “replace.” If the action is set to deny, any new IGMP join reports will be dropped. If the action is set to replace, the switch randomly removes an existing group and replaces it with the new multicast group.

Enabling IGMP Filtering and Throttling

Use the Multicast > IGMP Snooping > Filter (Configure General) page to enable IGMP filtering and throttling globally on the switch.

Parameters

These parameters are displayed:

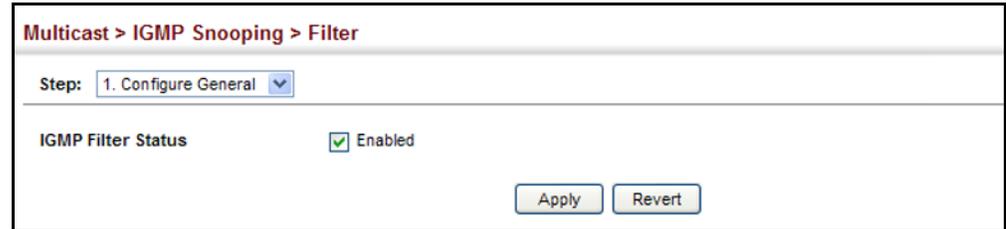
- ◆ **IGMP Filter Status** – Enables IGMP filtering and throttling globally for the switch. (Default: Disabled)

Web Interface

To enable IGMP filtering and throttling on the switch:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure General from the Step list.
3. Enable IGMP Filter Status.
4. Click Apply.

Figure 286: Enabling IGMP Filtering and Throttling



Configuring IGMP Filter Profiles

Use the Multicast > IGMP Snooping > Filter (Configure Profile – Add) page to create an IGMP profile and set its access mode. Then use the (Add Multicast Group Range) page to configure the multicast groups to filter.

Command Usage

Specify a range of multicast groups by entering a start and end IP address; or specify a single multicast group by entering the same IP address for the start and end of the range.

Parameters

These parameters are displayed:

Add

- ◆ **Profile ID** – Creates an IGMP profile. (Range: 1-4294967295)
- ◆ **Access Mode** – Sets the access mode of the profile; either permit or deny. (Default: Deny)

When the access mode is set to permit, IGMP join reports are processed when a multicast group falls within the controlled range. When the access mode is set to deny, IGMP join reports are only processed when the multicast group is not in the controlled range.

Add Multicast Group Range

- ◆ **Profile ID** – Selects an IGMP profile to configure.
- ◆ **Start Multicast IP Address** – Specifies the starting address of a range of multicast groups.
- ◆ **End Multicast IP Address** – Specifies the ending address of a range of multicast groups.

Web Interface

To create an IGMP filter profile and set its access mode:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure Profile from the Step list.
3. Select Add from the Action list.
4. Enter the number for a profile, and set its access mode.
5. Click Apply.

Figure 287: Creating an IGMP Filtering Profile

To show the IGMP filter profiles:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure Profile from the Step list.
3. Select Show from the Action list.

Figure 288: Showing the IGMP Filtering Profiles Created

	Profile ID	Action Mode
<input type="checkbox"/>	1	Permit
<input type="checkbox"/>	2	Deny
<input type="checkbox"/>	3	Deny
<input type="checkbox"/>	4294967295	Deny

To add a range of multicast groups to an IGMP filter profile:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure Profile from the Step list.
3. Select Add Multicast Group Range from the Action list.
4. Select the profile to configure, and add a multicast group address or range of addresses.
5. Click Apply.

Figure 289: Adding Multicast Groups to an IGMP Filtering Profile

The screenshot shows the configuration page for adding multicast groups. The breadcrumb is "Multicast > IGMP Snooping > Filter". The "Step" is set to "2. Configure Profile" and the "Action" is "Add Multicast Group Range". The "Profile ID" is set to "19". The "Start Multicast IP Address" is "239.2.3.1" and the "End Multicast IP Address" is "239.2.3.200". There are "Apply" and "Revert" buttons at the bottom right.

To show the multicast groups configured for an IGMP filter profile:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure Profile from the Step list.
3. Select Show Multicast Group Range from the Action list.
4. Select the profile for which to display this information.

Figure 290: Showing the Groups Assigned to an IGMP Filtering Profile

The screenshot shows the configuration page for showing multicast groups. The breadcrumb is "Multicast > IGMP Snooping > Filter". The "Step" is "2. Configure Profile" and the "Action" is "Show Multicast Group Range". The "Profile ID" is set to "1". Below the form is a table titled "Multicast IP Address Range List" with a "Total: 2" count. The table has three columns: a checkbox, "Start Multicast IP Address", and "End Multicast IP Address". There are two rows of data. At the bottom are "Delete" and "Revert" buttons.

<input type="checkbox"/>	Start Multicast IP Address	End Multicast IP Address
<input type="checkbox"/>	224.1.1.1	224.1.1.5
<input type="checkbox"/>	224.1.1.10	224.1.1.20

Configuring IGMP Filtering and Throttling for Interfaces Use the Multicast > IGMP Snooping > Filter (Configure Interface) page to assign and IGMP filter profile to interfaces on the switch, or to throttle multicast traffic by limiting the maximum number of multicast groups an interface can join at the same time.

Command Usage

- ◆ IGMP throttling sets a maximum number of multicast groups that a port can join at the same time. When the maximum number of groups is reached on a port, the switch can take one of two actions; either “deny” or “replace.” If the action is set to deny, any new IGMP join reports will be dropped. If the action is set to replace, the switch randomly removes an existing group and replaces it with the new multicast group.

Parameters

These parameters are displayed:

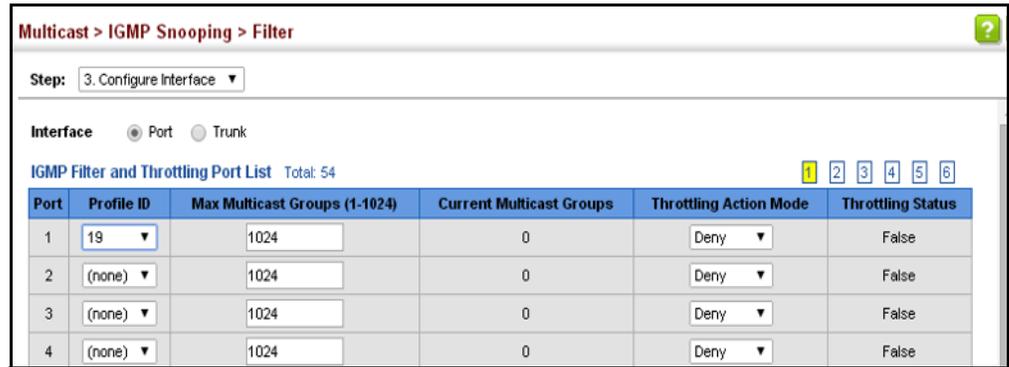
- ◆ **Interface** – Port or trunk identifier.
An IGMP profile or throttling setting can be applied to a port or trunk. When ports are configured as trunk members, the trunk uses the settings applied to the first port member in the trunk.
- ◆ **Profile ID** – Selects an existing profile to assign to an interface.
- ◆ **Max Multicast Groups** – Sets the maximum number of multicast groups an interface can join at the same time. (Range: 1-1024; Default: 1024)
- ◆ **Current Multicast Groups** – Displays the current multicast groups the interface has joined.
- ◆ **Throttling Action Mode** – Sets the action to take when the maximum number of multicast groups for the interface has been exceeded. (Default: Deny)
 - **Deny** - The new multicast group join report is dropped.
 - **Replace** - The new multicast group replaces an existing group.
- ◆ **Throttling Status** – Indicates if the throttling action has been implemented on the interface. (Options: True or False)

Web Interface

To configure IGMP filtering or throttling for a port or trunk:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure Interface from the Step list.
3. Select a profile to assign to an interface, then set the maximum number of allowed multicast groups and the throttling response.
4. Click Apply.

Figure 291: Configuring IGMP Filtering and Throttling Interface Settings



MLD Snooping (Snooping and Query for IPv6)

Multicast Listener Discovery (MLD) snooping operates on IPv6 traffic and performs a similar function to IGMP snooping for IPv4. That is, MLD snooping dynamically configures switch ports to limit IPv6 multicast traffic so that it is forwarded only to ports with users that want to receive it. This reduces the flooding of IPv6 multicast packets in the specified VLANs.

There are two versions of the MLD protocol, version 1 and version 2. MLDv1 control packets include Listener Query, Listener Report, and Listener Done messages (equivalent to IGMPv2 query, report, and leave messages). MLDv2 control packets include MLDv2 query and report messages, as well as MLDv1 report and done messages.

Remember that IGMP Snooping and MLD Snooping are independent functions, and can therefore both function at the same time.

Configuring MLD Snooping and Query Parameters

Use the Multicast > MLD Snooping > General page to configure the switch to forward multicast traffic intelligently. Based on the MLD query and report messages, the switch forwards multicast traffic only to the ports that request it. This prevents the switch from broadcasting the traffic to all ports and possibly disrupting network performance.

Parameters

These parameters are displayed:

- ◆ **MLD Snooping Status** – When enabled, the switch will monitor network traffic to determine which hosts want to receive multicast traffic. (Default: Disabled)
- ◆ **Querier Status** – When enabled, the switch can serve as the querier for MLDv2 snooping if elected. The querier is responsible for asking hosts if they want to receive multicast traffic. (Default: Disabled)

An IPv6 address must be configured on the VLAN interface from which the querier will act if elected. When serving as the querier, the switch uses this IPv6 address as the query source address.

The querier will not start or will disable itself after having started if it detects an IPv6 multicast router on the network.

- ◆ **Robustness** – MLD Snooping robustness variable. A port will be removed from the receiver list for a multicast service when no MLD reports are detected in response to a number of MLD queries. The robustness variable sets the number of queries on ports for which there is no report. (Range: 2-10 Default: 2)

- ◆ **Query Interval** – The interval between sending MLD general queries. (Range 60-125 seconds; Default: 125 seconds)

This attribute applies when the switch is serving as the querier.

An MLD general query message is sent by the switch at the interval specified by this attribute. When this message is received by downstream hosts, all receivers build an MLD report for the multicast groups they have joined.

- ◆ **Query Max Response Time** – The maximum response time advertised in MLD general queries. (Range: 5-25 seconds; Default: 10 seconds)

This attribute controls how long the host has to respond to an MLD Query message before the switch deletes the group if it is the last member.

- ◆ **Router Port Expiry Time** – The time the switch waits after the previous querier stops before it considers the router port (i.e., the interface that had been receiving query packets) to have expired. (Range: 300-500 seconds; Default: 300 seconds)

- ◆ **MLD Snooping Version** – The protocol version used for compatibility with other devices on the network. This is the MLD version the switch uses to send snooping reports. (Range: 1-2; Default: 2)

- ◆ **Unknown Multicast Mode** – The action for dealing with unknown multicast packets. Options include:

- **Flood** – Floods any received IPv6 multicast packets that have not been requested by a host to all ports in the VLAN.
- **To Router Port** – Forwards any received IPv6 multicast packets that have not been requested by a host to ports that are connected to a detected multicast router. (This is the default action.)

Web Interface

To configure general settings for MLD Snooping:

1. Click Multicast, MLD Snooping, General.
2. Adjust the settings as required.

3. Click Apply.

Figure 292: Configuring General Settings for MLD Snooping

Multicast > MLD Snooping > General

MLD Snooping Status Enabled

Querier Status Enabled

Robustness (2-10)

Query Interval (60-125) seconds

Query Max Response Time (5-25) seconds

Router Port Expiry Time (300-500) seconds

MLD Snooping Version (1-2)

Unknown Multicast Mode

Apply Revert

Setting Immediate Leave Status for MLD Snooping per Interface

Use the Multicast > MLD Snooping > Interface page to configure Immediate Leave status for a VLAN.

Parameters

These parameters are displayed:

- ◆ **VLAN** – A VLAN identification number. (Range: 1-4094)
- ◆ **Immediate Leave Status** – Immediately deletes a member port of an IPv6 multicast service when a leave packet is received at that port and immediate leave is enabled for the parent VLAN. (Default: Disabled)

If MLD immediate-leave is *not* used, a multicast router (or querier) will send a group-specific query message when an MLD group leave message is received. The router/querier stops forwarding traffic for that group only if no host replies to the query within the specified timeout period.

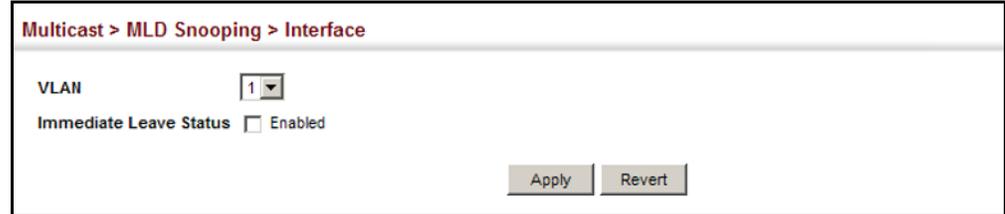
If MLD immediate-leave is enabled, the switch assumes that only one host is connected to the interface. Therefore, immediate leave should only be enabled on an interface if it is connected to only one MLD-enabled device, either a service host or a neighbor running MLD snooping.

Web Interface

To configure immediate leave for MLD Snooping:

1. Click Multicast, MLD Snooping, Interface.
2. Select a VLAN, and set the status for immediate leave.
3. Click Apply.

Figure 293: Configuring Immediate Leave for MLD Snooping



Multicast > MLD Snooping > Interface

VLAN

Immediate Leave Status Enabled

Apply Revert

Specifying Static Interfaces for an IPv6 Multicast Router

Use the Multicast > MLD Snooping > Multicast Router (Add Static Multicast Router) page to statically attach an interface to an IPv6 multicast router/switch.

Depending on your network connections, MLD snooping may not always be able to locate the MLD querier. Therefore, if the MLD querier is a known multicast router/switch connected over the network to an interface (port or trunk) on the switch, you can manually configure that interface to join all the current multicast groups.

Command Usage

MLD Snooping must be enabled globally on the switch (see [“Configuring MLD Snooping and Query Parameters” on page 462](#)) before a multicast router port can take effect.

Parameters

These parameters are displayed:

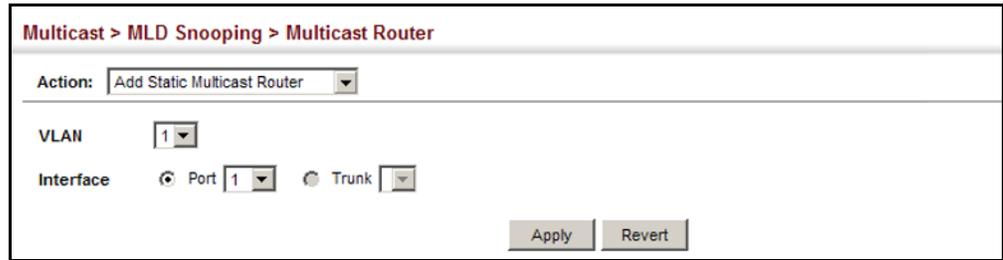
- ◆ **VLAN** – Selects the VLAN which is to propagate all IPv6 multicast traffic coming from the attached multicast router. (Range: 1-4094)
- ◆ **Interface** – Activates the Port or Trunk scroll down list.
- ◆ **Port or Trunk** – Specifies the interface attached to a multicast router.

Web Interface

To specify a static interface attached to a multicast router:

1. Click Multicast, MLD Snooping, Multicast Router.
2. Select Add Static Multicast Router from the Action list.
3. Select the VLAN which will forward all the corresponding IPv6 multicast traffic, and select the port or trunk attached to the multicast router.
4. Click Apply.

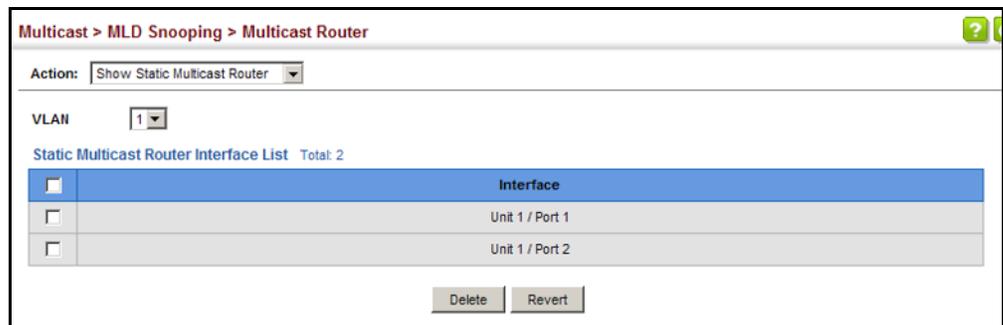
Figure 294: Configuring a Static Interface for an IPv6 Multicast Router



To show the static interfaces attached to a multicast router:

1. Click Multicast, MLD Snooping, Multicast Router.
2. Select Show Static Multicast Router from the Action list.
3. Select the VLAN for which to display this information.

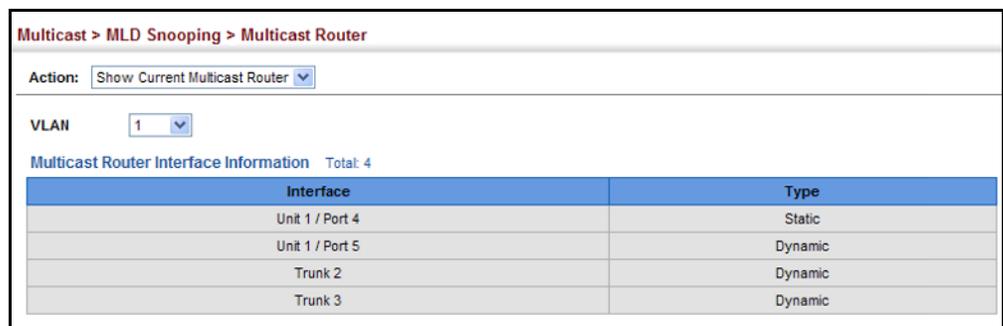
Figure 295: Showing Static Interfaces Attached an IPv6 Multicast Router



To show all the interfaces attached to a multicast router:

1. Click Multicast, MLD Snooping, Multicast Router.
2. Select Current Multicast Router from the Action list.
3. Select the VLAN for which to display this information. Ports in the selected VLAN which are attached to a neighboring multicast router/switch are displayed.

Figure 296: Showing Current Interfaces Attached an IPv6 Multicast Router



Assigning Interfaces to IPv6 Multicast Services

Use the Multicast > MLD Snooping > MLD Member (Add Static Member) page to statically assign an IPv6 multicast service to an interface.

Multicast filtering can be dynamically configured using MLD snooping and query messages (see [“Configuring MLD Snooping and Query Parameters” on page 462](#)). However, for certain applications that require tighter control, it may be necessary to statically configure a multicast service on the switch. First add all the ports attached to participating hosts to a common VLAN, and then assign the multicast service to that VLAN group.

Command Usage

- ◆ Static multicast addresses are never aged out.
- ◆ When a multicast address is assigned to an interface in a specific VLAN, the corresponding traffic can only be forwarded to ports within that VLAN.

Parameters

These parameters are displayed:

- ◆ **VLAN** – Specifies the VLAN which is to propagate the multicast service. (Range: 1-4094)
- ◆ **Multicast IPv6 Address** – The IP address for a specific multicast service.
- ◆ **Interface** – Activates the Port or Trunk scroll down list.
- ◆ **Port or Trunk** – Specifies the interface assigned to a multicast group.
- ◆ **Type** (Show Current Member) – Shows if this multicast stream was statically configured by the user, discovered by MLD Snooping, or is a data stream to which no other ports are subscribing (i.e., the stream is flooded onto VLAN instead of being trapped to the CPU for processing, or is being processed by MVR6).

Web Interface

To statically assign an interface to an IPv6 multicast service:

1. Click Multicast, MLD Snooping, MLD Member.
2. Select Add Static Member from the Action list.
3. Select the VLAN that will propagate the multicast service, specify the interface attached to a multicast service (through an MLD-enabled switch or multicast router), and enter the multicast IP address.
4. Click Apply.

Figure 297: Assigning an Interface to an IPv6 Multicast Service

Multicast > MLD Snooping > MLD Member

Action: Add Static Member

VLAN: 1

Multicast IPv6 Address: FF00:0:0:0:0:10C

Interface: Port 1

Apply Revert

To show the static interfaces assigned to an IPv6 multicast service:

1. Click Multicast, MLD Snooping, MLD Member.
2. Select Show Static Member from the Action list.
3. Select the VLAN for which to display this information.

Figure 298: Showing Static Interfaces Assigned to an IPv6 Multicast Service

Multicast > MLD Snooping > MLD Member

Action: Show Static Member

VLAN: 1

MLD Member Interface List Total: 8

	Multicast IPv6 Address	Interface
<input type="checkbox"/>	FF02::01:01:01:01	Unit 1 / Port 1
<input type="checkbox"/>	FF02::01:01:01:02	Unit 1 / Port 2
<input type="checkbox"/>	FF01::1	Unit 1 / Port 12
<input type="checkbox"/>	FF01::2	Unit 1 / Port 13
<input type="checkbox"/>	FF01::3	Unit 1 / Port 14
<input type="checkbox"/>	FF01::4	Unit 1 / Port 15
<input type="checkbox"/>	FF01::5	Unit 1 / Port 16
<input type="checkbox"/>	FF02::01:01:01:FF	Trunk 3

Delete Revert

To display information about all IPv6 multicast groups, MLD Snooping or multicast routing must first be enabled on the switch. To show all of the interfaces statically or dynamically assigned to an IPv6 multicast service:

1. Click Multicast, MLD Snooping, MLD Member.
2. Select Show Current Member from the Action list.
3. Select the VLAN for which to display this information.

Figure 299: Showing Current Interfaces Assigned to an IPv6 Multicast Service

Multicast IPv6 Address	Interface	Type
FF02::01:01:01:01	Unit 1 / Port 1	User
FF02::01:01:01:02	Unit 1 / Port 2	User
FF01::1	Unit 1 / Port 12	User
FF11::2	Unit 1 / Port 13	Multicast Data
FF11::3	Unit 1 / Port 14	User
FF11::4	Unit 1 / Port 15	User
FF11::5	Unit 1 / Port 16	User
FF02::01:01:01:FF	Trunk 3	User
FF03::01:01:01:FF	Trunk 5	MLD Snooping

Showing MLD Snooping Groups and Source List

Use the Multicast > MLD Snooping > Group Information page to display known multicast groups, member ports, the means by which each group was learned, and the corresponding source list.

Parameters

These parameters are displayed:

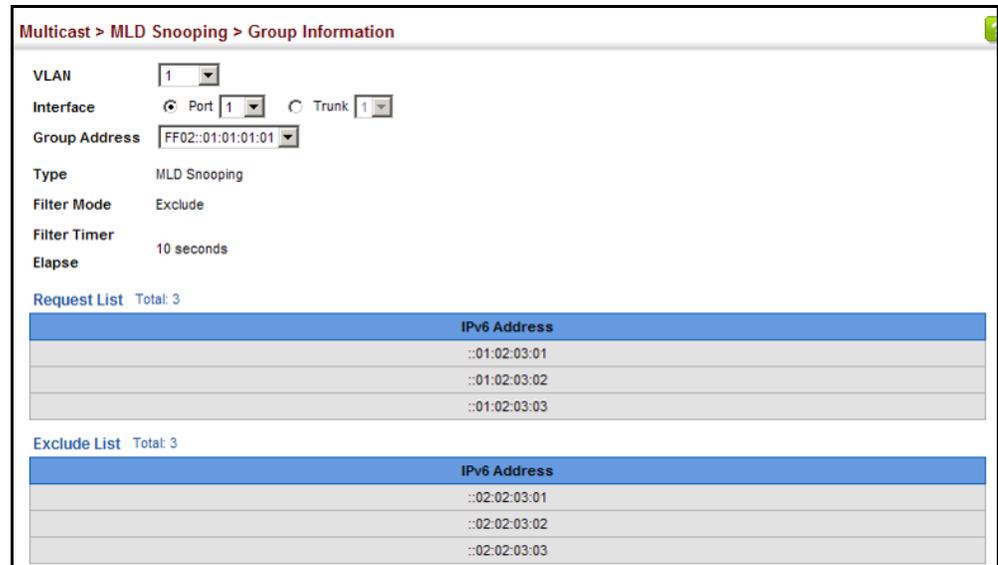
- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Interface** – Port or trunk identifier.
- ◆ **Group Address** – The IP address for a specific multicast service.
- ◆ **Type** – The means by which each group was learned – MLD Snooping or Multicast Data.
- ◆ **Filter Mode** – The filter mode is used to summarize the total listening state of a multicast address to a minimum set such that all nodes' listening states are respected. In Include mode, the router only uses the request list, indicating that the reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the hosts' source-list. In Exclude mode, the router uses both the request list and exclude list, indicating that the reception of packets sent to the given multicast address is requested from all IP source addresses, except for those listed in the exclude source-list and for any other sources where the source timer status has expired.
- ◆ **Filter Timer Elapse** – The Filter timer is only used when a specific multicast address is in Exclude mode. It represents the time for the multicast address filter mode to expire and change to Include mode.
- ◆ **Request List** – Sources included on the router's request list.
- ◆ **Exclude List** – Sources included on the router's exclude list.

Web Interface

To display known MLD multicast groups:

1. Click Multicast, MLD Snooping, Group Information.
2. Select the port or trunk, and then select a multicast service assigned to that interface.

Figure 300: Showing IPv6 Multicast Services and Corresponding Sources



Layer 3 IGMP (Query used with Multicast Routing)

IGMP Snooping – IGMP Snooping (page 437) is a key part of the overall set of functions required to support multicast filtering. It is used to passively monitor IGMP service requests from multicast clients, and dynamically configure the switch ports which need to forward multicast traffic.

IGMP Query – Multicast query is used to poll each known multicast group for active members, and dynamically configure the switch ports which need to forward multicast traffic. Layer 3 IGMP Query, as described below, is used in conjunction with both Layer 2 IGMP Snooping and multicast routing.

IGMP – This protocol includes a form of multicast query specifically designed to work with multicast routing. A router periodically asks its hosts if they want to receive multicast traffic. It then propagates service requests on to any upstream multicast router to ensure that it will continue to receive the multicast service. IGMP can be enabled for individual VLAN interfaces (page 474).



Note: Multicast Routing Discovery (MRD) is used to discover which interfaces are attached to multicast routers. (For a description of this protocol, see “Multicast Router Discovery” on [page 445](#).)

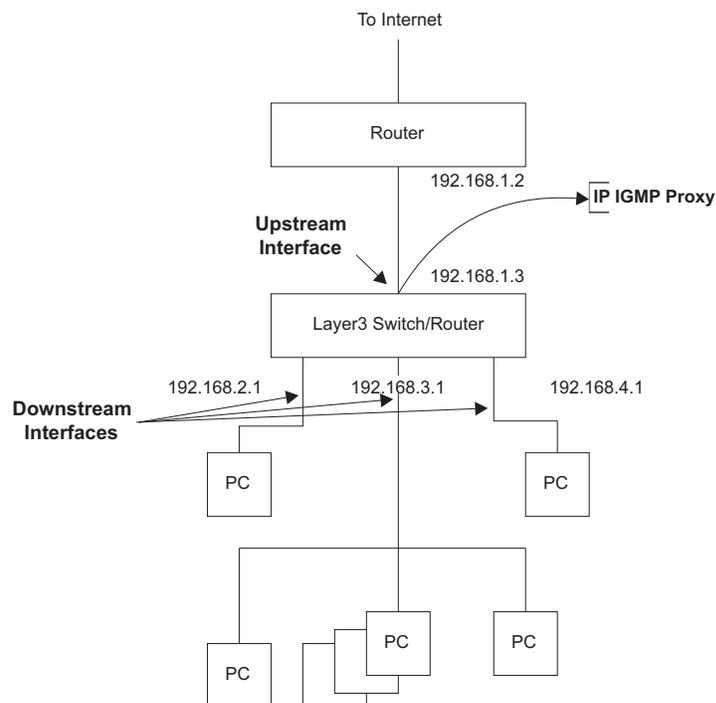
IGMP Proxy – A device can learn about the multicast service requirements of hosts attached to its downstream interfaces, proxy this group membership information to the upstream router, and forward multicast packets based on that information.

Configuring IGMP Proxy Routing

Use the Multicast > IGMP > Proxy page to configure IGMP Proxy Routing.

In simple network topologies, it is sufficient for a device to learn multicast requirements from its downstream interfaces and proxy this group membership information to the upstream router. Multicast packets can then be forwarded downstream based solely upon that information. This mechanism, known as IGMP proxy routing, enables the system to issue IGMP host messages on behalf of hosts that the system has discovered through standard IGMP interfaces.

Figure 301: IGMP Proxy Routing



Using IGMP proxy routing to forward multicast traffic on edge switches greatly reduces the processing load on those devices by not having to run more complicated multicast routing protocols such as PIM. It also makes the proxy devices independent of the multicast routing protocols used by core routers.

IGMP proxy routing uses a tree topology, where the root of the tree is connected to a complete multicast infrastructure (with the upstream interface connected to the Internet as shown in the figure above). In such a simple topology, it is sufficient to

send the group membership information learned upstream, and then to forward multicast packets based upon that information to the downstream hosts. For the switch, IGMP proxy routing has only one upstream connection to the core network side and multiple downstream connections to the customer side.

The IGMP proxy routing tree must be manually configured by designating one upstream interface and multiple downstream interfaces on each proxy device. No other multicast routers except for the proxy devices can exist within the tree, and the root of the tree must be connected to a wider multicast infrastructure. Note that this protocol is limited to a single administrative domain.

In more complicated scenarios where the topology is not a tree (such as when there are diverse paths to multiple sources), a more robust failover mechanism should be used. If more than one administrative domain is involved, a multicast routing protocol should be used instead of IGMP proxy.

To enable IGMP proxy service, follow these steps:

1. Enable IP multicasting globally on the router (see [“Configuring Global Settings for Multicast Routing”](#) on page 602).
2. Enable IGMP on the downstream interfaces which require proxy multicast service (see [“Configuring IGMP Interface Parameters”](#) on page 474).
3. Enable IGMP proxy on the interface that is attached to an upstream multicast router using the proxy settings described in this section.
4. Optional – Indicate how often the system will send unsolicited reports to the upstream router using the Multicast > IGMP > Proxy page as described later in this section.

Command Usage

- ◆ When IGMP proxy is enabled on an interface, that interface is known as the upstream or host interface. This interface performs only the host portion of IGMP by sending IGMP membership reports, and automatically disables IGMP router functions.
- ◆ Interfaces with IGMP enabled, but not located in the direction of the multicast tree root are known as downstream or router interfaces. These interfaces perform the standard IGMP router functions by maintaining a database of all IGMP subscriptions on the downstream interface. IGMP must therefore be enabled on all interfaces which require proxy multicast service.
- ◆ The system periodically checks the multicast route table for (*,G) any-source multicast forwarding entries. When changes occur in the downstream IGMP groups, an IGMP state change report is created and sent to the upstream router.
- ◆ If there is an IGMPv1 or IGMPv2 querier on the upstream network, then the proxy device will act as an IGMPv1 or IGMPv2 host on the upstream interface

accordingly, and set the v1/v2 query present timer to indicate that there is an active v1/v2 querier in this VLAN. Otherwise, it will act as an IGMPv3 host.

- ◆ Multicast routing protocols are not supported when IGMP proxy service is enabled.
- ◆ Only one upstream interface is supported on the system.
- ◆ A maximum of 1024 multicast entries are supported.

Parameters

These parameters are displayed:

- ◆ **VLAN** – VLAN interface on which to configure IGMP proxy service. (Range: 1-4094)
- ◆ **IGMP Proxy Status** – Enables IGMP proxy service for multicast routing, forwarding IGMP membership information monitored on downstream interfaces onto the upstream interface in a summarized report. (Default: Disabled)
- ◆ **Unsolicited Report Interval** – Specifies how often the upstream interface should transmit unsolicited IGMP reports. (Range: 1-65535 seconds; Default: 400 seconds)

Web Interface

To configure IGMP Proxy Routing:

1. Click Multicast, IGMP, Proxy.
2. Select the upstream interface, enable the IGMP Proxy Status, and modify the interval for unsolicited IGMP reports if required.
3. Click Apply.

Figure 302: Configuring IGMP Proxy Routing

Multicast > IGMP > Proxy

VLAN

IGMP Proxy Status Enabled

Unsolicited Report Interval (1-65535) seconds

Configuring IGMP Interface Parameters

Use the Multicast > IGMP > Interface page to configure interface settings for IGMP.

The switch uses IGMP (Internet Group Management Protocol) to query for any attached hosts that want to receive a specific multicast service. The hosts may respond with several types of IP multicast messages. Hosts respond to queries with report messages that indicate which groups they want to join or the groups to which they already belong. If a router does not receive a report message within a specified period of time, it will prune that interface from the multicast tree. A host can also submit a join message at any time without waiting for a query from the router. Hosts can also signal when they no longer want to receive traffic for a specific group by sending a leave-group message.

If more than one router on the LAN is performing IP multicasting, one of these is elected as the “querier” and assumes the role of querying for group members. It then propagates the service request up to any neighboring multicast router to ensure that it will continue to receive the multicast service. The parameters described in this section are used to control Layer 3 IGMP and query functions.



Note: IGMP Protocol Status should be enabled on all the interfaces that need to support downstream multicast hosts (as described in this section).

Note: IGMP is disabled when multicast routing is disabled (see [“Enabling Multicast Routing Globally” on page 602](#)).

Parameters

These parameters are displayed:

- ◆ **VLAN** – VLAN interface bound to a primary IP address. (Range: 1-4093)
- ◆ **IGMP Protocol Status** – Enables IGMP (including IGMP query functions) on a VLAN interface. (Default: Disabled)

When a multicast routing protocol, such as PIM, is enabled, IGMP is also enabled.
- ◆ **IGMP Version** – Configures the IGMP version used on an interface. (Options: Version 1-3; Default: Version 2)
- ◆ **Robustness Variable** – Specifies the robustness (or expected packet loss) for this interface. The robustness value is used in calculating the appropriate range for other IGMP variables, such as the Group Membership Interval, as well as the Other Querier Present Interval, and the Startup Query Count (RFC 2236). (Range: 1-255; Default: 2)

Routers adopt the robustness value from the most recently received query. If the querier's robustness variable (QRV) is zero, indicating that the QRV field does not contain a declared robustness value, the switch will set the robustness variable to the value statically configured by this command. If the QRV exceeds 7, the maximum value of the QRV field, the robustness value is set to zero,

meaning that this device will not advertise a QRV in any query messages it subsequently sends.

- ◆ **Query Interval** – Configures the frequency at which host query messages are sent. (Range: 1-255; Default: 125 seconds)

Multicast routers send host query messages to determine the interfaces that are connected to downstream hosts requesting a specific multicast service. Only the designated multicast router for a subnet sends host query messages, which are addressed to the multicast address 224.0.0.1, and use a time-to-live (TTL) value of 1.

For IGMP Version 1, the designated router is elected according to the multicast routing protocol that runs on the LAN. But for IGMP Version 2 and 3, the designated querier is the lowest IP-addressed multicast router on the subnet.

- ◆ **Query Max Response Time** – Configures the maximum response time advertised in IGMP queries. (Range: 0-255 tenths of a second; Default: 10 seconds)

IGMPv1 does not support a configurable maximum response time for query messages. It is fixed at 10 seconds for IGMPv1.

By varying the Query Maximum Response Time, the burstiness of IGMP messages passed on the subnet can be tuned; where larger values make the traffic less bursty, as host responses are spread out over a larger interval.

The number of seconds represented by the maximum response interval must be less than the Query Interval.

- ◆ **Last Member Query Interval** – The frequency at which to send IGMP group-specific or IGMPv3 group-source-specific query messages in response to receiving a group-specific or group-source-specific leave message. (Range: 1-255 tenths of a second; Default: 1 second)

When the switch receives an IGMPv2 or IGMPv3 leave message from a host that wants to leave a multicast group, source or channel, it sends a number of group-specific or group-source-specific query messages as defined by the Last Member Query Count at intervals defined by the Last Member Query Interval. If no response is received after this period, the switch stops forwarding for the group, source or channel.

- ◆ **Querier** – Device currently serving as the IGMP querier for this multicast service. A querier can only be displayed if IGMP multicasting is enabled, the VLAN for this entry is up, and is configured with a valid IP address.

Web Interface

To configure IGMP interface settings:

1. Click Multicast, IGMP, Interface.
2. Select each interface that will support IGMP (Layer 3), and set the required IGMP parameters.

3. Click Apply.

Figure 303: Configuring IGMP Interface Settings

Multicast > IGMP > Interface

VLAN	1
IGMP Protocol Status	<input checked="" type="checkbox"/> Enabled
IGMP Version (1-3)	2
Robustness Variable (1-255)	1
Query Interval (1-255)	125 seconds
Query Max Response Time (0-255)	100 * 0.1 seconds
Last Member Query Interval (1-255)	10 * 0.1 seconds
Querier	192.168.1.254

Apply Revert

Configuring Static IGMP Group Membership

Use the Multicast > IGMP > Static Group page to manually propagate traffic from specific multicast groups onto the specified VLAN interface.

Command Usage

- ◆ Group addresses within the entire multicast group address range can be specified. However, if any address within the source-specific multicast (SSM) address range (default 232/8) is specified, but no source address is included, the request to join the multicast group will fail unless the next node up the reverse path tree has statically mapped this group to a specific source address. Also, if an address outside of the SSM address range is specified, and a specific source address is included in the command, the request to join the multicast group will also fail if the next node up the reverse path tree has enabled the PIM-SSM protocol.
- ◆ If a static group is configured for an any-source multicast (*,G), a source address cannot subsequently be defined for this group without first deleting the entry.
- ◆ If a static group is configured for one or more source-specific multicasts (S,G), an any-source multicast (*,G) cannot subsequently be defined for this group without first deleting all of the associated (S,G) entries.
- ◆ The switch supports a maximum of 64 static group entries.

Parameters

These parameters are displayed:

- ◆ **VLAN** – VLAN interface to assign as a static member of the specified multicast group. (Range: 1-4094)
- ◆ **Static Group Address** – An IP multicast group address. (The group addresses specified cannot be in the range of 224.0.0.1 - 239.255.255.255.)

- ◆ **Source Address** – The source address of a multicast server transmitting traffic to the specified multicast group address.

Web Interface

To configure static IGMP groups:

1. Click Multicast, IGMP, Static Group.
2. Select Add from the Action list.
3. Select a VLAN interface to be assigned as a static multicast group member, and then specify the multicast group. If source-specific multicasting is supported by the next hop router in the reverse path tree for the specified multicast group, then the source address should also be specified.
4. Click Apply.

Figure 304: Configuring Static IGMP Groups

Multicast > IGMP > Static Group

Action: Add

VLAN: 1

Static Group Address: 239.2.3.1

Source Address: 192.168.1.200 (optional)

Apply Revert

To display configured static IGMP groups:

1. Click Multicast, IGMP, Static Group.
2. Select Show from the Action list.
3. Click Apply.

Figure 305: Showing Static IGMP Groups

Multicast > IGMP > Static Group

Action: Show

VLAN: 1

IGMP Static Group List Total: 1

	Static Group Address	Source Address
<input type="checkbox"/>	239.2.3.1	192.168.1.200

Apply Revert

Displaying Multicast Group Information

When IGMP (Layer 3) is enabled on the switch, use the Multicast > IGMP > Group Information pages to display the current multicast groups learned through IGMP. When IGMP (Layer 3) is disabled and IGMP (Layer 2) is enabled, the active multicast groups can be viewed on the Multicast > IGMP Snooping > Forwarding Entry page (see [page 452](#)).

Command Usage

To display information about multicast groups, IGMP must first be enabled on the interface to which a group has been assigned (see [“Configuring IGMP Interface Parameters” on page 474](#)), and multicast routing must be enabled globally on the system (see [“Configuring Global Settings for Multicast Routing” on page 602](#)).

Parameters

These parameters are displayed:

Show Information

- ◆ **VLAN** – VLAN identifier. The selected entry must be a configured IP interface. (Range: 1-4094)
- ◆ **Group Address** – IP multicast group address with subscribers directly attached or downstream from the switch.
- ◆ **Last Reporter** – The IP address of the source of the last membership report received for this multicast group address on this interface.
- ◆ **Up Time** – The time elapsed since this entry was created. (Depending on the elapsed time, information may displayed for w:weeks, d:days, h:hours, m:minutes, or s:seconds.)
- ◆ **Expire** – The time remaining before this entry will be aged out. (Default: 260 seconds)

This parameter displays “stopped” if the Group Mode is INCLUDE.

- ◆ **V1 Timer** – The time remaining until the switch assumes that there are no longer any IGMP Version 1 members on the IP subnet attached to this interface.
 - If the switch receives an IGMP Version 1 Membership Report, it sets a timer to note that there are Version 1 hosts present which are members of the group for which it heard the report.
 - If there are Version 1 hosts present for a particular group, the switch will ignore any Leave Group messages that it receives for that group.

Show Details

The following additional information is displayed on this page:

- ◆ **VLAN** – VLAN identifier. The selected entry must be a configured IP interface. (Range: 1-4094)
- ◆ **Group Address** – IP multicast group address with subscribers directly attached or downstream from the switch, or a static multicast group assigned to this interface.
- ◆ **Interface** – The interface on the switch that has received traffic directed to the multicast group address.
- ◆ **Up Time** – The time elapsed since this entry was created. (Depending on the elapsed time, information may displayed for w:weeks, d:days, h:hours, m:minutes, or s:seconds.)
- ◆ **Group Mode** – In INCLUDE mode, reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the source-list parameter. In EXCLUDE mode, reception of packets sent to the given multicast address is requested from all IP source addresses, except for those listed in the source-list parameter and for any other sources where the source timer status has expired.
- ◆ **Group Source List** – A list of zero or more IP unicast addresses from which multicast reception is desired or not desired, depending on the filter mode.
 - **Source Address** – The address of one of the multicast servers transmitting traffic to the specified group.
 - **Up Time** – The time elapsed since this entry was created. (Depending on the elapsed time, information may displayed for w:weeks, d:days, h:hours, m:minutes, or s:seconds.)
 - **V3 Expire** – The time remaining before this entry will be aged out. The V3 label indicates that the expire time is only provided for sources learned through IGMP Version 3. (The default is 260 seconds.)
 - **Forward** – Indicates whether or not traffic will be forwarded from the multicast source.

Web Interface

To display the current multicast groups learned through IGMP:

1. Click Multicast, IGMP, Group Information.
2. Select Show Information from the Action list.
3. Select a VLAN. The selected entry must be a configured IP interface.

Figure 306: Displaying Multicast Groups Learned from IGMP (Information)

Multicast > IGMP > Group Information

Action: Show Information

VLAN: 1

IGMP Group Information List Total: 1

Group Address	Last Reporter	Up Time	Expire	V1 Timer
224.0.17.17	192.168.1.0	0:00:01	0:04:19	0:00:00

To display detailed information about the current multicast groups learned through IGMP:

1. Click Multicast, IGMP, Group Information.
2. Select Show Details from the Action list.
3. Select a VLAN. The selected entry must be a configured IP interface.

Figure 307: Displaying Multicast Groups Learned from IGMP

Multicast > IGMP > Group Information

Action: Show Details

VLAN: 1

Group Address: 224.1.1.1

Interface: VLAN 1

Up Time: 0h:12m:42s

Group Mode: Exclude

Last Reporter: 0.0.0.0

Group Source List Total: 3

Source Address	Up Time	V3 Expire	Forward
10.2.2.2	0h:1m:7s	10h:20m:0s	YES
11.2.2.2	0h:1m:7s	10h:20m:0s	YES
12.2.2.2	0h:1m:7s	10h:20m:0s	NO

IP Configuration

This chapter describes how to configure an IP interface for management access to the switch over the network. This switch supports both IP Version 4 and Version 6, and can be managed simultaneously through either of these address types. You can manually configure a specific IPv4 or IPv6 address, or direct the switch to obtain an IPv4 address from a DHCP server when it is powered on. An IPv6 global unicast or link-local address can be manually configured, or a link-local address can be dynamically generated.

This chapter provides information on network functions including:

- ◆ [IPv4 Configuration](#) – Sets an IPv4 address for management access.
- ◆ [IPv6 Configuration](#) – Sets an IPv6 address for management access.

Setting the Switch's IP Address (IP Version 4)

This section describes how to configure an initial IPv4 interface for management access over the network, or for creating an interface to multiple subnets. This switch supports both IPv4 and IPv6, and can be managed through either of these address types. For information on configuring the switch with an IPv6 address, see [“Setting the Switch's IP Address \(IP Version 6\)” on page 485](#).

Use the IP > General > Routing Interface (Add Address) page to configure an IPv4 address for the switch. An IPv4 address is obtained via DHCP by default for VLAN 1. To configure a static address, you need to change the switch's default settings to values that are compatible with your network. You may also need to establish a default gateway between the switch and management stations that exist on another network segment (if no routing protocols are enabled).

You can direct the device to obtain an address from a DHCP server, or manually configure a static IP address. Valid IP addresses consist of four decimal numbers, 0 to 255, separated by periods. Anything other than this format will not be accepted.

Command Usage

- ◆ This section describes how to configure a single local interface for initial access to the switch. To configure multiple IP interfaces, set up an IP interface for each VLAN.
- ◆ Once an IP address has been assigned to an interface, routing between different interfaces on the switch is enabled.

- ◆ To enable routing between interfaces defined on this switch and external network interfaces, you must configure static routes ([page 526](#)) or use dynamic routing; i.e., RIP ([page 544](#)), OSPFv2 ([page 562](#)), OSPFv3, or BGPv4. Note that OSPFv3 and BGPv4 are only supported through the Command Line Interface.
- ◆ The precedence for configuring IP interfaces is the IP > General > Routing Interface (Add) menu, static routes ([page 526](#)), and then dynamic routing.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of the configured VLAN (1-4094). By default, all ports on the switch are members of VLAN 1. However, the management station can be attached to a port belonging to any VLAN, as long as that VLAN has been assigned an IP address.
- ◆ **IP Address Mode** – Specifies whether IP functionality is enabled via manual configuration (User Specified), or Dynamic Host Configuration Protocol (DHCP), or Boot Protocol (BOOTP). If DHCP/BOOTP is enabled, IP will not function until a reply has been received from the server. Requests will be broadcast periodically by the switch for an IP address. DHCP/BOOTP responses can include the IP address, subnet mask, and default gateway. (Default: DHCP)

- ◆ **IP Address Type** – Specifies a primary or secondary IP address. An interface can have only one primary IP address, but can have many secondary IP addresses. In other words, secondary addresses need to be specified if more than one IP subnet can be accessed through this interface. For initial configuration, set this parameter to Primary. (Options: Primary, Secondary; Default: Primary)

Note that a secondary address cannot be configured prior to setting the primary IP address, and the primary address cannot be removed if a secondary address is still present. Also, if any router or switch in a network segment uses a secondary address, all other routers/switches in that segment must also use a secondary address from the same network or subnet address space.

- ◆ **IP Address** – IP Address of the VLAN. Valid IP addresses consist of four numbers, 0 to 255, separated by periods. (Default: None)



Note: You can manage the switch through any configured IP interface.

- ◆ **Subnet Mask** – This mask identifies the host address bits used for routing to specific subnets. (Default: None)
- ◆ **Restart DHCP** – Requests a new IP address from the DHCP server.

Web Interface

To set a static IPv4 address for the switch:

1. Click IP, General, Routing Interface.
2. Select Add Address from the Action list.
3. Select any configured VLAN, set IP Address Mode to "User Specified," set IP Address Type to "Primary" if no address has yet been configured for this interface, and then enter the IP address and subnet mask.
4. Click Apply.

Figure 308: Configuring a Static IPv4 Address

IP > General > Routing Interface

Action: Add Address

VLAN: 1

IP Address Mode: User Specified

IP Address Type: Primary

IP Address: 192.168.0.2

Subnet Mask: 255.255.255.0

Restart DHCP [Click this button to restart DHCP service.](#)

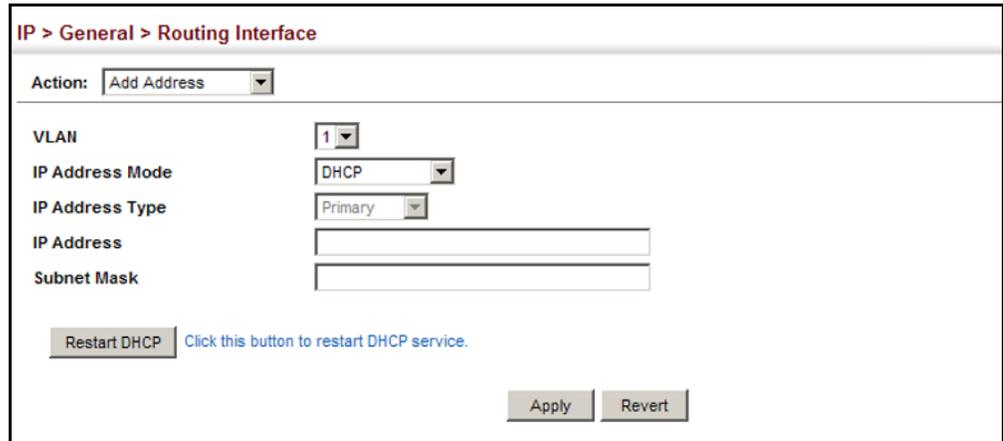
Apply Revert

To obtain a dynamic IPv4 address through DHCP/BOOTP for the switch:

1. Click IP, General, Routing Interface.
2. Select Add Address from the Action list.
3. Select any configured VLAN, and set IP Address Mode to "BOOTP" or "DHCP."
4. Click Apply to save your changes.
5. Then click Restart DHCP to immediately request a new address.

IP will be enabled but will not function until a BOOTP or DHCP reply is received. Requests are broadcast every few minutes using exponential backoff until IP configuration information is obtained from a BOOTP or DHCP server.

Figure 309: Configuring a Dynamic IPv4 Address



The screenshot shows a web interface for configuring a dynamic IPv4 address. The breadcrumb navigation is "IP > General > Routing Interface". The "Action" dropdown is set to "Add Address". The "VLAN" dropdown is set to "1". The "IP Address Mode" dropdown is set to "DHCP". The "IP Address Type" dropdown is set to "Primary". There are empty input fields for "IP Address" and "Subnet Mask". A "Restart DHCP" button is present with a link "Click this button to restart DHCP service." At the bottom right, there are "Apply" and "Revert" buttons.



Note: The switch will also broadcast a request for IP configuration settings on each power reset.

Note: If you lose the management connection, make a console connection to the switch and enter "show ip interface" to determine the new switch address.

Renewing DHCP – DHCP may lease addresses to clients indefinitely or for a specific period of time. If the address expires or the switch is moved to another network segment, you will lose management access to the switch. In this case, you can reboot the switch or submit a client request to restart DHCP service via the CLI.

If the address assigned by DHCP is no longer functioning, you will not be able to renew the IP settings via the web interface. You can only restart DHCP service via the web interface if the current address is still available.

To show the IPv4 address configured for an interface:

1. Click IP, General, Routing Interface.
2. Select Show Address from the Action list.
3. Select an entry from the VLAN list.

Figure 310: Showing the IPv4 Address Configured for an Interface



Setting the Switch's IP Address (IP Version 6)

This section describes how to configure an initial IPv6 interface for management access over the network, or for creating an interface to multiple subnets. This switch supports both IPv4 and IPv6, and can be managed through either of these address types. For information on configuring the switch with an IPv4 address, see [“Setting the Switch's IP Address \(IP Version 4\)” on page 481](#).

Command Usage

- ◆ IPv6 includes two distinct address types – link-local unicast and global unicast. A link-local address makes the switch accessible over IPv6 for all devices attached to the same local subnet. Management traffic using this kind of address cannot be passed by any router outside of the subnet. A link-local address is easy to set up, and may be useful for simple networks or basic troubleshooting tasks. However, to connect to a larger network with multiple segments, the switch must be configured with a global unicast address. A link-local address can be dynamically assigned (using the Configure Interface page) or manually configured (using the Add IPv6 Address page). A global unicast address can only be manually configured (using the Add IPv6 Address page).
- ◆ An IPv6 global unicast or link-local address can be manually configured (using the Add IPv6 Address page), or a link-local address can be dynamically generated (using the Configure Interface page).

Configuring the IPv6 Default Gateway

Use the IP > IPv6 Configuration (Configure Global) page to configure an IPv6 default gateway for the switch.

Parameters

These parameters are displayed:

- ◆ **Default Gateway** – Sets the IPv6 address of the default next hop router to use when no routing information is known about an IPv6 address.
 - If no routing protocol is enabled nor static route defined, you must define a gateway if the target device is located in a different subnet.

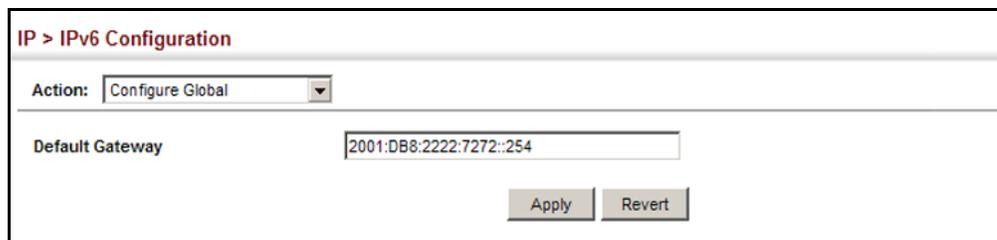
- If a routing protocol is enabled ([page 543](#)), you can still define a static route ([page 526](#)) to ensure that traffic to the designated address or subnet passes through a preferred gateway.
- An IPv6 default gateway can only be successfully set when a network interface that directly connects to the gateway has been configured on the switch.
- An IPv6 address must be configured according to RFC 2373 "IPv6 Addressing Architecture," using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields.

Web Interface

To configure an IPv6 default gateway for the switch:

1. Click IP, IPv6 Configuration.
2. Select Configure Global from the Action list.
3. Enter the IPv6 default gateway.
4. Click Apply.

Figure 311: Configuring the IPv6 Default Gateway



The screenshot shows a web interface for configuring IPv6 settings. At the top, it says "IP > IPv6 Configuration". Below that, there is a dropdown menu for "Action" with "Configure Global" selected. Underneath, there is a text input field for "Default Gateway" containing the address "2001:DB8:2222:7272::254". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

Configuring IPv6 Interface Settings

Use the IP > IPv6 Configuration (Configure Interface) page to configure general IPv6 settings for the selected VLAN, including explicit configuration of a link local interface address, the MTU size, and neighbor discovery protocol settings for duplicate address detection and the neighbor solicitation interval.

Command Usage

- ◆ The switch must be configured with a link-local address. The option to explicitly enable IPv6 creates a link-local address, but will not generate a global IPv6 address. The global unicast address must be manually configured (see ["Configuring an IPv6 Address" on page 491](#)).
- ◆ IPv6 Neighbor Discovery Protocol supersedes IPv4 Address Resolution Protocol in IPv6 networks. IPv6 nodes on the same network segment use Neighbor Discovery to discover each other's presence, to determine each other's link-layer addresses, to find routers and to maintain reachability information about the paths to active neighbors. The key parameters used to facilitate this process

are the number of attempts made to verify whether or not a duplicate address exists on the same network segment, and the interval between neighbor solicitations used to verify reachability information.

Parameters

These parameters are displayed:

VLAN Mode

- ◆ **VLAN** – ID of a configured VLAN which is to be used for management access, or as a standard interface for a subnet. By default, all ports on the switch are members of VLAN 1. However, the management station can be attached to a port belonging to any VLAN, as long as that VLAN has been assigned an IP address. (Range: 1-4094)
- ◆ **Address Autoconfig** – Enables stateless autoconfiguration of an IPv6 address on an interface and enables IPv6 functionality on that interface. The network portion of the address is based on prefixes received in IPv6 router advertisement messages, and the host portion is automatically generated using the modified EUI-64 form of the interface identifier (i.e., the switch's MAC address).
 - If a link local address has not yet been assigned to this interface, this command will dynamically generate one. The link local address is made with an address prefix in the range of FE80~FEBF and a host portion based the switch's MAC address in modified EUI-64 format. It will also generate a global unicast address if a global prefix is included in received router advertisements.
 - When DHCPv6 is started, the switch may attempt to acquire an IP address prefix through stateful address autoconfiguration. If router advertisements have the "other stateful configuration" flag set, the switch will attempt to acquire other non-address configuration information (such as a default gateway).
 - If auto-configuration is not selected, then an address must be manually configured using the Add IPv6 Address page described below.
- ◆ **Enable IPv6 Explicitly** – Enables IPv6 on an interface and assigns it a link-local address. Note that when an explicit address is assigned to an interface, IPv6 is automatically enabled, and cannot be disabled until all assigned addresses have been removed. (Default: Disabled)

Disabling this parameter does not disable IPv6 for an interface that has been explicitly configured with an IPv6 address.

- ◆ **MTU** – Sets the size of the maximum transmission unit (MTU) for IPv6 packets sent on an interface. (Range: 1280-65535 bytes; Default: 1500 bytes)
 - The maximum value set in this field cannot exceed the MTU of the physical interface for lower layer packets, which is currently fixed at 1500 bytes.
 - If a non-default value is configured, an MTU option is included in the router advertisements sent from this device. This option is provided to ensure that

all nodes on a link use the same MTU value in cases where the link MTU is not otherwise well known.

- IPv6 routers do not fragment IPv6 packets forwarded from other routers. However, traffic originating from an end-station connected to an IPv6 router may be fragmented.
- All devices on the same physical medium must use the same MTU in order to operate correctly.
- IPv6 must be enabled on an interface before the MTU can be set. If an IPv6 address has not been assigned to the switch, "N/A" is displayed in the MTU field.

◆ **ND DAD Attempts** – The number of consecutive neighbor solicitation messages sent on an interface during duplicate address detection. (Range: 0-600, Default: 1)

- Configuring a value of 0 disables duplicate address detection.
- Duplicate address detection determines if a new unicast IPv6 address already exists on the network before it is assigned to an interface.
- Duplicate address detection is stopped on any interface that has been suspended (see ["Configuring VLAN Groups" on page 147](#)). While an interface is suspended, all unicast IPv6 addresses assigned to that interface are placed in a "pending" state. Duplicate address detection is automatically restarted when the interface is administratively re-activated.
- An interface that is re-activated restarts duplicate address detection for all unicast IPv6 addresses on the interface. While duplicate address detection is performed on the interface's link-local address, the other IPv6 addresses remain in a "tentative" state. If no duplicate link-local address is found, duplicate address detection is started for the remaining IPv6 addresses.
- If a duplicate address is detected, it is set to "duplicate" state, and a warning message is sent to the console. If a duplicate link-local address is detected, IPv6 processes are disabled on the interface. If a duplicate global unicast address is detected, it is not used. All configuration commands associated with a duplicate address remain configured while the address is in "duplicate" state.
- If the link-local address for an interface is changed, duplicate address detection is performed on the new link-local address, but not for any of the IPv6 global unicast addresses already associated with the interface.

◆ **ND NS Interval** – The interval between transmitting IPv6 neighbor solicitation messages on an interface. (Range: 1000-3600000 milliseconds;

Default: 1000 milliseconds is used for neighbor discovery operations, 0 milliseconds is advertised in router advertisements.

This attribute specifies the interval between transmitting neighbor solicitation messages when resolving an address, or when probing the reachability of a neighbor. Therefore, avoid using very short intervals for normal IPv6 operations.

When a non-default value is configured, the specified interval is used both for router advertisements and by the router itself.

- ◆ **ND Reachable-Time** – The amount of time that a remote IPv6 node is considered reachable after some reachability confirmation event has occurred. (Range: 0-3600000 milliseconds)

Default: 30000 milliseconds is used for neighbor discovery operations, 0 milliseconds is advertised in router advertisements.

- The time limit configured by this parameter allows the router to detect unavailable neighbors. During the neighbor discover process, an IPv6 node will multicast neighbor solicitation messages to search for neighbor nodes. For a neighbor node to be considered reachable, it must respond to the neighbor soliciting node with a neighbor advertisement message to become a confirmed neighbor, after which the reachable timer will be considered in effect for subsequent unicast IPv6 layer communications.
- This time limit is included in all router advertisements sent out through an interface, ensuring that nodes on the same link use the same time value.
- Setting the time limit to 0 means that the configured time is unspecified by this router.

RA Guard Mode

- ◆ **Interface** – Shows port or trunk configuration page.
- ◆ **RA Guard** – Blocks incoming Router Advertisement and Router Redirect packets. (Default: Disabled)

IPv6 Router Advertisements (RA) convey information that enables nodes to auto-configure on the network. This information may include the default router address taken from the observed source address of the RA message, as well as on-link prefix information. However, note that unintended misconfigurations, or possibly malicious attacks on the network, may lead to bogus RAs being sent, which in turn can cause operational problems for hosts on the network.

RA Guard can be used to block RAs and Router Redirect (RR) messages on the specified interface. Determine which interfaces are connected to known routers, and enable RA Guard on all other untrusted interfaces.

Web Interface

To configure general IPv6 settings for the switch:

1. Click IP, IPv6 Configuration.
2. Select Configure Interface from the Action list.
3. Select VLAN mode.
4. Specify the VLAN to configure.

5. Enable IPv6 Explicitly to automatically configure a link-local address and enable IPv6 on the selected interface. (To manually configure the link-local address, use the Add IPv6 Address page.) Set the MTU size, the maximum number of duplicate address detection messages, the neighbor solicitation message interval, and the amount of time that a remote IPv6 node is considered reachable.
6. Click Apply.

Figure 312: Configuring General Settings for an IPv6 Interface

The screenshot shows the 'IP > IPv6 Configuration' page. At the top, there is a breadcrumb 'IP > IPv6 Configuration' and an 'Action:' dropdown menu set to 'Configure Interface'. Below this, several configuration options are listed:

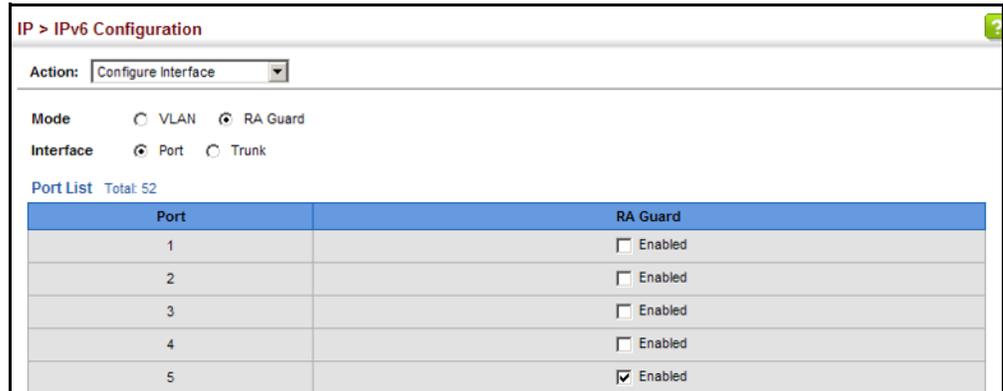
- VLAN:** A dropdown menu showing '1'.
- Address Autoconfig:** A checkbox labeled 'Enabled' which is currently unchecked.
- Enable IPv6 Explicitly:** A checkbox labeled 'Enabled' which is currently unchecked.
- MTU (1280-65535):** A text input field containing '1500' followed by the unit 'bytes'.
- ND DAD Attempts (0-600):** A text input field containing '1'.
- ND HS Interval (1000-3600000):** A text input field containing '1000' followed by the unit 'ms'.
- ND Reachable-Time (0-3600000):** A text input field containing '30000' followed by the unit 'ms'.

At the bottom of the configuration area, there is a button labeled 'Restart DHCPv6' and a link that says 'Click this button to restart DHCPv6 service.' At the very bottom of the page, there are two buttons: 'Apply' and 'Revert'.

To configure RA Guard for the switch:

1. Click IP, IPv6 Configuration.
2. Select Configure Interface from the Action list.
3. Select RA Guard mode.
4. Enable RA Guard for untrusted interfaces.
5. Click Apply.

Figure 313: Configuring RA Guard for an IPv6 Interface



Configuring an IPv6 Address Use the IP > IPv6 Configuration (Add IPv6 Address) page to configure an initial IPv6 interface for management access over the network, or for creating an interface to multiple subnets.

Command Usage

- ◆ All IPv6 addresses must be formatted according to RFC 2373 “IPv6 Addressing Architecture,” using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields.
- ◆ The switch must always be configured with a link-local address. Therefore explicitly enabling IPv6 (see [“Configuring IPv6 Interface Settings” on page 486](#)) or manually assigning a global unicast address will also automatically generate a link-local unicast address. The prefix length for a link-local address is fixed at 64 bits, and the host portion of the default address is based on the modified EUI-64 (Extended Universal Identifier) form of the interface identifier (i.e., the physical MAC address). Alternatively, you can manually configure the link-local address by entering the full address with a network prefix in the range of FE80~FEBF.
- ◆ To connect to a larger network with multiple subnets, you must configure a global unicast address. There are several alternatives to configuring this address type:
 - It can be manually configured by specifying the entire network prefix and prefix length, and using the EUI-64 form of the interface identifier to automatically create the low-order 64 bits in the host portion of the address.
 - You can also manually configure the global unicast address by entering the full address and prefix length.
- ◆ You can configure multiple IPv6 global unicast addresses per interface, but only one link-local address per interface.

- ◆ If a duplicate link-local address is detected on the local segment, this interface is disabled and a warning message displayed on the console. If a duplicate global unicast address is detected on the network, the address is disabled on this interface and a warning message displayed on the console.
- ◆ When an explicit address is assigned to an interface, IPv6 is automatically enabled, and cannot be disabled until all assigned addresses have been removed.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of a configured VLAN which is to be used for management access, or for creating an interface to multiple subnets. By default, all ports on the switch are members of VLAN 1. However, the management station can be attached to a port belonging to any VLAN, as long as that VLAN has been assigned an IP address. (Range: 1-4094)
- ◆ **Address Type** – Defines the address type configured for this interface.
 - **Global** – Configures an IPv6 global unicast address with a full IPv6 address including the network prefix and host address bits, followed by a forward slash, and a decimal value indicating how many contiguous bits (from the left) of the address comprise the prefix (i.e., the network portion of the address).
 - **EUI-64** (Extended Universal Identifier) – Configures an IPv6 address for an interface using an EUI-64 interface ID in the low order 64 bits.
 - When using EUI-64 format for the low-order 64 bits in the host portion of the address, the value entered in the IPv6 Address field includes the network portion of the address, and the prefix length indicates how many contiguous bits (starting at the left) of the address comprise the prefix (i.e., the network portion of the address). Note that the value specified in the IPv6 Address field may include some of the high-order host bits if the specified prefix length is less than 64 bits. If the specified prefix length exceeds 64 bits, then the bits used in the network portion of the address will take precedence over the interface identifier.
 - IPv6 addresses are 16 bytes long, of which the bottom 8 bytes typically form a unique host identifier based on the device's MAC address. The EUI-64 specification is designed for devices that use an extended 8-byte MAC address. For devices that still use a 6-byte MAC address (also known as EUI-48 format), it must be converted into EUI-64 format by inverting the universal/local bit in the address and inserting the hexadecimal number FFFE between the upper and lower three bytes of the MAC address.

For example, if a device had an EUI-48 address of 28-9F-18-1C-82-35, the global/local bit must first be inverted to meet EUI-64 requirements (i.e., 1 for globally defined addresses and 0 for locally defined

addresses), changing 28 to 2A. Then the two bytes FFFE are inserted between the OUI (i.e., organizationally unique identifier, or company identifier) and the rest of the address, resulting in a modified EUI-64 interface identifier of 2A-9F-18-FF-FE-1C-82-35.

- This host addressing method allows the same interface identifier to be used on multiple IP interfaces of a single device, as long as those interfaces are attached to different subnets.
- **Link Local** – Configures an IPv6 link-local address.
 - The address prefix must be in the range of FE80~FEBF.
 - You can configure only one link-local address per interface.
 - The specified address replaces a link-local address that was automatically generated for the interface.
- ◆ **IPv6 Address** – IPv6 address assigned to this interface.

Web Interface

To configure an IPv6 address:

1. Click IP, IPv6 Configuration.
2. Select Add IPv6 Address from the Action list.
3. Specify the VLAN to configure, select the address type, and then enter an IPv6 address and prefix length.
4. Click Apply.

Figure 314: Configuring an IPv6 Address

The screenshot shows the 'IP > IPv6 Configuration' web interface. At the top, there is a breadcrumb trail 'IP > IPv6 Configuration'. Below it, the 'Action:' dropdown menu is set to 'Add IPv6 Address'. The configuration fields are as follows:

VLAN	1
Address Type	Global
IPv6 Address	2001:DB8:2222:7272::72/96

At the bottom right of the form, there are two buttons: 'Apply' and 'Revert'.

Showing IPv6 Addresses Use the IP > IPv6 Configuration (Show IPv6 Address) page to display the IPv6 addresses assigned to an interface.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of a configured VLAN. By default, all ports on the switch are members of VLAN 1. However, the management station can be attached to a port belonging to any VLAN, as long as that VLAN has been assigned an IP address. (Range: 1-4094)
- ◆ **IPv6 Address Type** – The address type (Global, EUI-64, Link Local).
- ◆ **IPv6 Address** – An IPv6 address assigned to this interface.

In addition to the unicast addresses assigned to an interface, a host is also required to listen to the all-nodes multicast addresses FF01::1 (interface-local scope) and FF02::1 (link-local scope).

FF01::1/16 is the transient interface-local multicast address for all attached IPv6 nodes, and FF02::1/16 is the link-local multicast address for all attached IPv6 nodes. The interface-local multicast address is only used for loopback transmission of multicast traffic. Link-local multicast addresses cover the same types as used by link-local unicast addresses, including all nodes (FF02::1), all routers (FF02::2), and solicited nodes (FF02::1:FFXX:XXXX) as described below.

A node is also required to compute and join the associated solicited-node multicast addresses for every unicast and anycast address it is assigned. IPv6 addresses that differ only in the high-order bits, e.g. due to multiple high-order prefixes associated with different aggregations, will map to the same solicited-node address, thereby reducing the number of multicast addresses a node must join. In this example, FF02::1:FF90:0/104 is the solicited-node multicast address which is formed by taking the low-order 24 bits of the address and appending those bits to the prefix.

Note that the solicited-node multicast address (link-local scope FF02) is used to resolve the MAC addresses for neighbor nodes since IPv6 does not support the broadcast method used by the Address Resolution Protocol in IPv4.

These additional addresses are displayed by the CLI (see the “show ip interface” command in the *CLI Reference Guide*).

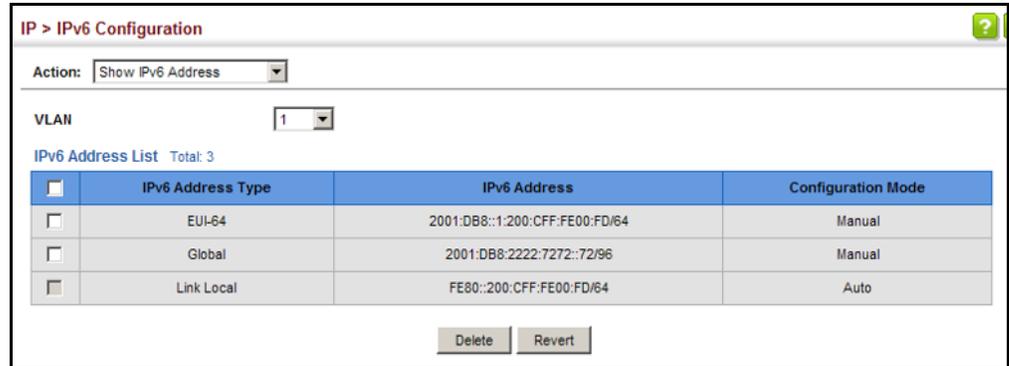
- ◆ **Configuration Mode** – Indicates if this address was automatically generated or manually configured.

Web Interface

To show the configured IPv6 addresses:

1. Click IP, IPv6 Configuration.
2. Select Show IPv6 Address from the Action list.
3. Select a VLAN from the list.

Figure 315: Showing Configured IPv6 Addresses



Showing the IPv6 Neighbor Cache Use the IP > IPv6 Configuration (Show IPv6 Neighbor Cache) page to display the IPv6 addresses detected for neighbor devices.

Parameters

These parameters are displayed:

Table 36: Show IPv6 Neighbors - display description

Field	Description
IPv6 Address	IPv6 address of neighbor
Age	The time since the address was verified as reachable (in seconds). A static entry is indicated by the value "Permanent."
Link-layer Address	Physical layer MAC address.
State	The following states are used for dynamic entries: <ul style="list-style-type: none"> ◆ Incomplete - Address resolution is being carried out on the entry. A neighbor solicitation message has been sent to the multicast address of the target, but it has not yet returned a neighbor advertisement message. ◆ Invalid - An invalidated mapping. Setting the state to invalid dis-associates the interface identified with this entry from the indicated mapping (RFC 4293). ◆ Reachable - Positive confirmation was received within the last ReachableTime interval that the forward path to the neighbor was functioning. While in REACH state, the device takes no special action when sending packets. ◆ Stale - More than the ReachableTime interval has elapsed since the last positive confirmation was received that the forward path was functioning. While in STALE state, the device takes no action until a packet is sent.

Table 36: Show IPv6 Neighbors - display description (Continued)

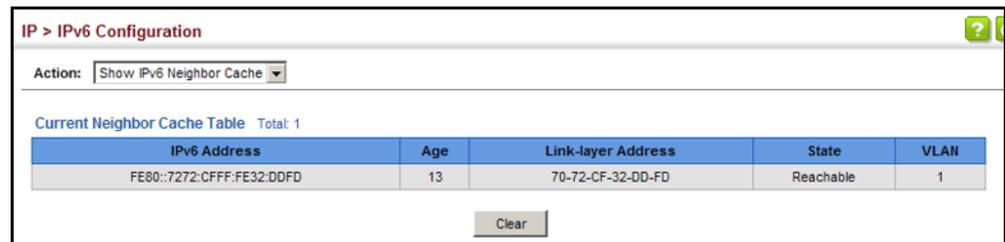
Field	Description
	<ul style="list-style-type: none"> ◆ Delay - More than the ReachableTime interval has elapsed since the last positive confirmation was received that the forward path was functioning. A packet was sent within the last DELAY_FIRST_PROBE_TIME interval. If no reachability confirmation is received within this interval after entering the DELAY state, the switch will send a neighbor solicitation message and change the state to PROBE. ◆ Probe - A reachability confirmation is actively sought by re-sending neighbor solicitation messages every RetransTimer interval until confirmation of reachability is received. ◆ Unknown - Unknown state. <p>The following states are used for static entries:</p> <ul style="list-style-type: none"> ◆ Incomplete -The interface for this entry is down. ◆ Permanent - Indicates a static entry. ◆ Reachable - The interface for this entry is up. Reachability detection is not applied to static entries in the IPv6 neighbor discovery cache.
VLAN	VLAN interface from which the address was reached.

Web Interface

To show neighboring IPv6 devices:

1. Click IP, IPv6 Configuration.
2. Select Show IPv6 Neighbors from the Action list.

Figure 316: Showing IPv6 Neighbors



Showing IPv6 Statistics Use the IP > IPv6 Configuration (Show Statistics) page to display statistics about IPv6 traffic passing through this switch.

Command Usage

This switch provides statistics for the following traffic types:

- ◆ **IPv6** – The Internet Protocol for Version 6 addresses provides a mechanism for transmitting blocks of data (often called packets or frames) from a source to a destination, where these network devices (that is, hosts) are identified by fixed length addresses. The Internet Protocol also provides for fragmentation and reassembly of long packets, if necessary, for transmission through “small packet” networks.

- ◆ **ICMPv6** – Internet Control Message Protocol for Version 6 addresses is a network layer protocol that transmits message packets to report errors in processing IPv6 packets. ICMP is therefore an integral part of the Internet Protocol. ICMP messages may be used to report various situations, such as when a datagram cannot reach its destination, when the gateway does not have the buffering capacity to forward a datagram, and when the gateway can direct the host to send traffic on a shorter route. ICMP is also used by routers to feed back information about more suitable routes (that is, the next hop router) to use for a specific destination.
- ◆ **UDP** – User Datagram Protocol provides a datagram mode of packet switched communications. It uses IP as the underlying transport mechanism, providing access to IP-like services. UDP packets are delivered just like IP packets – connection-less datagrams that may be discarded before reaching their targets. UDP is useful when TCP would be too complex, too slow, or just unnecessary.

Parameters

These parameters are displayed:

Table 37: Show IPv6 Statistics - display description

Field	Description
IPv6 Statistics	
<i>IPv6 Received</i>	
Total	The total number of input datagrams received by the interface, including those received in error.
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, IPv6 options, etc.
Too Big Errors	The number of input datagrams that could not be forwarded because their size exceeded the link MTU of outgoing interface.
No Routes	The number of input datagrams discarded because no route could be found to transmit them to their destination.
Address Errors	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 (e.g., ::0) and unsupported addresses (e.g., addresses with unallocated prefixes). For entities which are not IPv6 routers and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
Unknown Protocols	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 necessarily the input interface for some of the datagrams.
Truncated Packets	The number of input datagrams discarded because datagram frame didn't carry enough data.
Discards	The number of input IPv6 datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.

Table 37: Show IPv6 Statistics - display description (Continued)

Field	Description
Delivers	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 necessarily the input interface for some of the datagrams.
Reassembly Request Datagrams	The number of IPv6 fragments received which needed to be reassembled at this interface. Note that this counter is incremented at the interface to which these fragments were addressed which might not be necessarily the input interface for some of the fragments.
Reassembly Succeeded	The number of IPv6 datagrams successfully reassembled. Note that 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.
Reassembly Failed	The number of failures detected by the IPv6 re-assembly algorithm (for whatever reason: timed out, errors, etc.). Note that 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 necessarily the input interface for some of the fragments.
<i>IPv6 Transmitted</i>	
Forwards Datagrams	The number of output datagrams which this entity received and forwarded to their final destinations. In entities which do not act as IPv6 routers, this counter will include only those packets which were Source-Routed via this entity, and the Source-Route processing was successful. Note that for a successfully forwarded datagram the counter of the outgoing interface is incremented."
Requests	The total number of IPv6 datagrams which local IPv6 user-protocols (including ICMP) supplied to IPv6 in requests for transmission. Note that this counter does not include any datagrams counted in <code>ipv6IfStatsOutForwDatagrams</code> .
Discards	The number of output IPv6 datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (e.g., for lack of buffer space). Note that this counter would include datagrams counted in <code>ipv6IfStatsOutForwDatagrams</code> if any such packets met this (discretionary) discard criterion.
No Routes	The number of input datagrams discarded because no route could be found to transmit them to their destination.
Generated Fragments	The number of output datagram fragments that have been generated as a result of fragmentation at this output interface.
Fragment Succeeded	The number of IPv6 datagrams that have been successfully fragmented at this output interface.
Fragment Failed	The number of IPv6 datagrams that have been discarded because they needed to be fragmented at this output interface but could not be.
ICMPv6 Statistics	
<i>ICMPv6 received</i>	
Input	The total number of ICMP messages received by the interface which includes all those counted by <code>ipv6IfIcmpInErrors</code> . Note that this interface is the interface to which the ICMP messages were addressed which may not be necessarily the input interface for the messages.
Errors	The number of ICMP messages which the interface received but determined as having ICMP-specific errors (bad ICMP check sums, bad length, etc.).

Table 37: Show IPv6 Statistics - display description (Continued)

Field	Description
Destination Unreachable Messages	The number of ICMP Destination Unreachable messages received by the interface.
Packet Too Big Messages	The number of ICMP Packet Too Big messages received by the interface.
Time Exceeded Messages	The number of ICMP Time Exceeded messages received by the interface.
Parameter Problem Messages	The number of ICMP Parameter Problem messages received by the interface.
Echo Request Messages	The number of ICMP Echo (request) messages received by the interface.
Echo Reply Messages	The number of ICMP Echo Reply messages received by the interface.
Router Solicit Messages	The number of ICMP Router Solicit messages received by the interface.
Router Advertisement Messages	The number of ICMP Router Advertisement messages received by the interface.
Neighbor Solicit Messages	The number of ICMP Neighbor Solicit messages received by the interface.
Neighbor Advertisement Messages	The number of ICMP Neighbor Advertisement messages received by the interface.
Redirect Messages	The number of Redirect messages received by the interface.
Group Membership Query Messages	The number of ICMPv6 Group Membership Query messages received by the interface.
Group Membership Response Messages	The number of ICMPv6 Group Membership Response messages received by the interface.
Group Membership Reduction Messages	The number of ICMPv6 Group Membership Reduction messages received by the interface.
Multicast Listener Discovery Version 2 Reports	The number of MLDv2 reports received by the interface.
<i>ICMPv6 Transmitted</i>	
Output	The total number of ICMP messages which this interface attempted to send. Note that this counter includes all those counted by icmpOutErrors.
Destination Unreachable Messages	The number of ICMP Destination Unreachable messages sent by the interface.
Packet Too Big Messages	The number of ICMP Packet Too Big messages sent by the interface.
Time Exceeded Messages	The number of ICMP Time Exceeded messages sent by the interface.
Parameter Problem Message	The number of ICMP Parameter Problem messages sent by the interface.
Echo Request Messages	The number of ICMP Echo (request) messages sent by the interface.
Echo Reply Messages	The number of ICMP Echo Reply messages sent by the interface.
Router Solicit Messages	The number of ICMP Router Solicitation messages sent by the interface.
Router Advertisement Messages	The number of ICMP Router Advertisement messages sent by the interface.
Neighbor Solicit Messages	The number of ICMP Neighbor Solicit messages sent by the interface.

Table 37: Show IPv6 Statistics - display description (Continued)

Field	Description
Neighbor Advertisement Messages	The number of ICMP Router Advertisement messages sent by the interface.
Redirect Messages	The number of Redirect messages sent. For a host, this object will always be zero, since hosts do not send redirects.
Group Membership Query Messages	The number of ICMPv6 Group Membership Query messages sent by the interface.
Group Membership Response Messages	The number of ICMPv6 Group Membership Response messages sent.
Group Membership Reduction Messages	The number of ICMPv6 Group Membership Reduction messages sent.
Multicast Listener Discovery Version 2 Reports	The number of MLDv2 reports sent by the interface.
UDP Statistics	
Input	The total number of UDP datagrams delivered to UDP users.
No Port Errors	The total number of received UDP datagrams for which there was no application at the destination port.
Other Errors	The number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port.
Output	The total number of UDP datagrams sent from this entity.

Web Interface

To show the IPv6 statistics:

1. Click IP, IPv6 Configuration.
2. Select Show Statistics from the Action list.
3. Click IPv6, ICMPv6 or UDP.

Figure 317: Showing IPv6 Statistics (IPv6)

The screenshot shows a web interface for IP > IPv6 configuration. At the top, there is a breadcrumb 'IP > IPv6'. Below it, an 'Action:' dropdown menu is set to 'Show Statistics'. Underneath, there are three radio buttons for 'Type': 'IPv6' (selected), 'ICMPv6', and 'UDP'. The main content area is titled 'IPv6 Statistics' and contains a table of statistics. The table has two columns of statistics and a central column for values. A 'Clear' button is located at the bottom right of the statistics area.

IPv6 Statistics		
Total Received	55	Received Reassembled Succeeded
Received Header Errors	0	Received Reassembled Failed
Received Too Big Errors	0	Transmitted Forwards Datagrams
Received No Routes	0	Transmitted Requests
Received Address Errors	0	Transmitted Discards
Received Unknown Protocols	0	Transmitted No Routes
Received Truncated Packets	0	Transmitted Generated Fragments
Received Discards	0	Transmitted Fragment Succeeded
Received Delivers	55	Transmitted Fragment Failed
Received Reassembly Request Datagrams	0	

Figure 318: Showing IPv6 Statistics (ICMPv6)

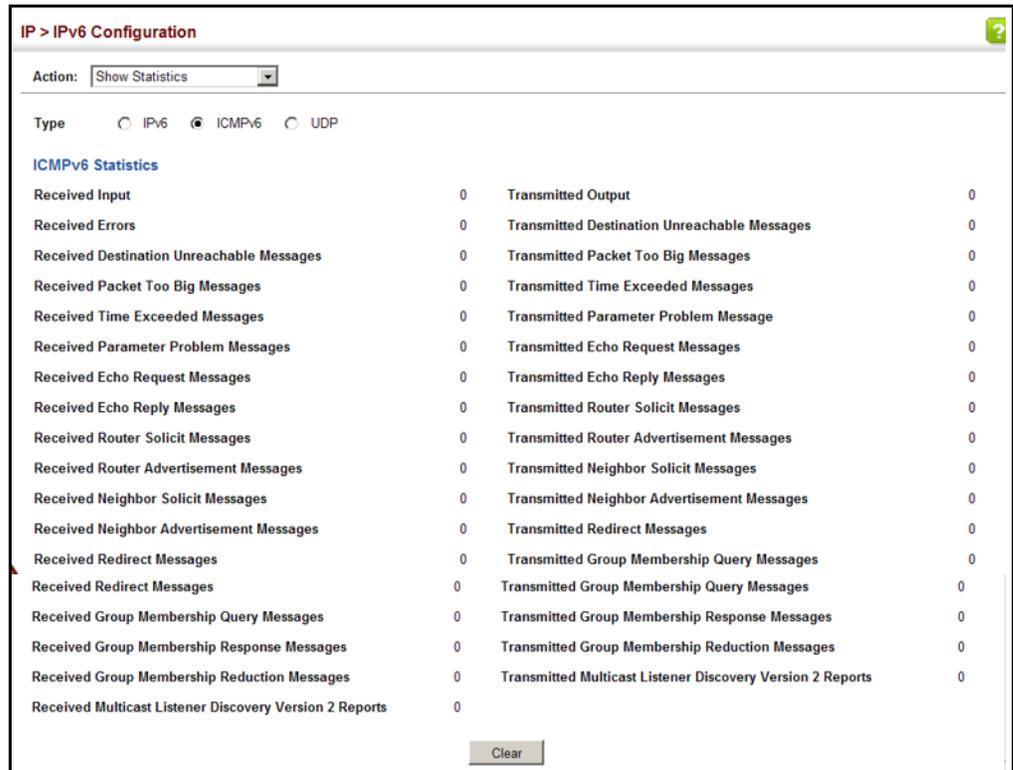
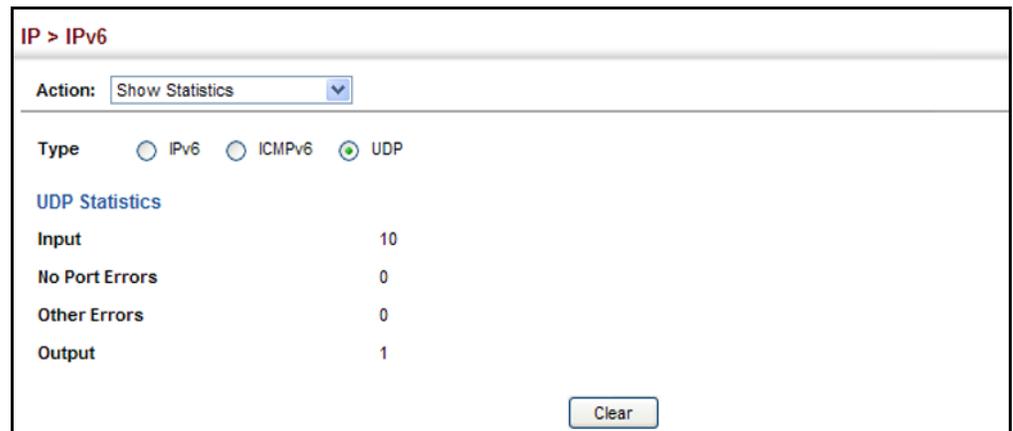


Figure 319: Showing IPv6 Statistics (UDP)



Showing the MTU for Responding Destinations Use the IP > IPv6 Configuration (Show MTU) page to display the maximum transmission unit (MTU) cache for destinations that have returned an ICMP packet-too-big message along with an acceptable MTU to this switch.

Parameters

These parameters are displayed:

Table 38: Show MTU - display description

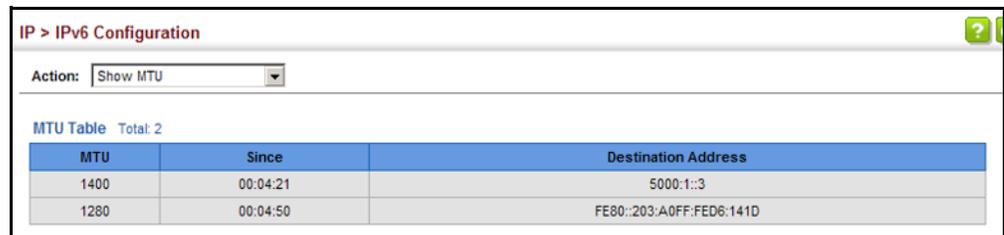
Field	Description
MTU	Adjusted MTU contained in the ICMP packet-too-big message returned from this destination, and now used for all traffic sent along this path.
Since	Time since an ICMP packet-too-big message was received from this destination.
Destination Address	Address which sent an ICMP packet-too-big message.

Web Interface

To show the MTU reported from other devices:

1. Click IP, IPv6 Configuration.
2. Select Show MTU from the Action list.

Figure 320: Showing Reported MTU Values



IP Services

This chapter describes the following IP services:

- ◆ **DNS** – Configures default domain names, identifies servers to use for dynamic lookup, and shows how to configure static entries.
- ◆ **DHCP Client** – Specifies the DHCP client identifier for an interface.
- ◆ **DHCP Relay** – Enables DHCP relay service, and defines the servers to which client requests are forwarded.



Note: For information on DHCP snooping which is included in this folder, see [“DHCP Snooping” on page 318](#).

Domain Name Service

DNS service on this switch allows host names to be mapped to IP addresses using static table entries or by redirection to other name servers on the network. When a client device designates this switch as a DNS server, the client will attempt to resolve host names into IP addresses by forwarding DNS queries to the switch, and waiting for a response.

You can manually configure entries in the DNS table used for mapping domain names to IP addresses, configure default domain names, or specify one or more name servers to use for domain name to address translation.

Configuring General DNS Service Parameters

Use the IP Service > DNS - General (Configure Global) page to enable domain lookup and set the default domain name.

Command Usage

- ◆ To enable DNS service on this switch, enable domain lookup status, and configure one or more name servers (see [“Configuring a List of Name Servers” on page 508](#)).
- ◆ If one or more name servers are configured, but DNS is not yet enabled and the switch receives a DHCP packet containing a DNS field with a list of DNS servers, then the switch will automatically enable DNS host name-to-address translation.

Parameters

These parameters are displayed:

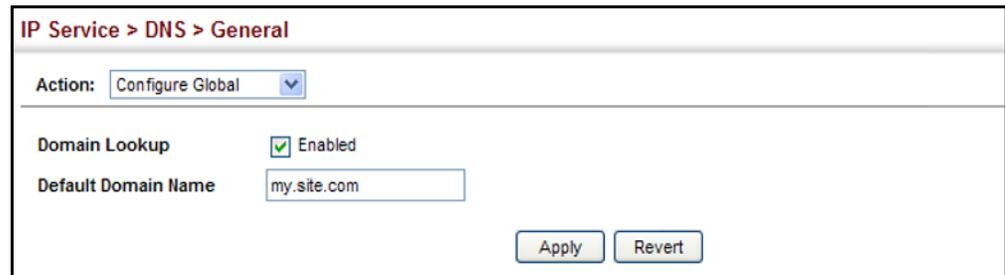
- ◆ **Domain Lookup** – Enables DNS host name-to-address translation. (Default: Disabled)
- ◆ **Default Domain Name** – Defines the default domain name appended to incomplete host names. Do not include the initial dot that separates the host name from the domain name. (Range: 1-127 alphanumeric characters)

Web Interface

To configure general settings for DNS:

1. Click IP Service, DNS.
2. Select Configure Global from the Action list.
3. Enable domain lookup, and set the default domain name.
4. Click Apply.

Figure 321: Configuring General Settings for DNS



IP Service > DNS > General

Action:

Domain Lookup Enabled

Default Domain Name

Configuring a List of Domain Names Use the IP Service > DNS - General (Add Domain Name) page to configure a list of domain names to be tried in sequential order.

Command Usage

- ◆ Use this page to define a list of domain names that can be appended to incomplete host names (i.e., host names passed from a client that are not formatted with dotted notation).
- ◆ If there is no domain list, the default domain name is used (see [“Configuring General DNS Service Parameters” on page 505](#)). If there is a domain list, the system will search it for a corresponding entry. If none is found, it will use the default domain name.
- ◆ When an incomplete host name is received by the DNS service on this switch and a domain name list has been specified, the switch will work through the domain list, appending each domain name in the list to the host name, and

checking with the specified name servers for a match (see “Configuring a List of Name Servers” on page 508).

Parameters

These parameters are displayed:

Domain Name – Name of the host. Do not include the initial dot that separates the host name from the domain name. (Range: 1-68 characters)

Web Interface

To create a list domain names:

1. Click IP Service, DNS.
2. Select Add Domain Name from the Action list.
3. Enter one domain name at a time.
4. Click Apply.

Figure 322: Configuring a List of Domain Names for DNS

The screenshot shows a web interface with the breadcrumb "IP Service > DNS > General". Below the breadcrumb, there is an "Action:" dropdown menu set to "Add Domain Name". Underneath, there is a "Domain Name" input field containing the text "sample.com.uk". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

To show the list domain names:

1. Click IP Service, DNS.
2. Select Show Domain Names from the Action list.

Figure 323: Showing the List of Domain Names for DNS

The screenshot shows the same web interface as Figure 322, but the "Action:" dropdown menu is now set to "Show Domain Names". Below the dropdown, there is a section titled "Domain Name List Total: 2". This section contains a table with two rows of domain names. Each row has a checkbox in the first column and the domain name in the second column. At the bottom right, there are two buttons: "Delete" and "Revert".

Domain Name
<input type="checkbox"/> google.com
<input type="checkbox"/> hinet.net

Configuring a List of Name Servers Use the IP Service > DNS - General (Add Name Server) page to configure a list of name servers to be tried in sequential order.

Command Usage

- ◆ To enable DNS service on this switch, configure one or more name servers, and enable domain lookup status (see [“Configuring General DNS Service Parameters” on page 505](#)).
- ◆ When more than one name server is specified, the servers are queried in the specified sequence until a response is received, or the end of the list is reached with no response.
- ◆ If all name servers are deleted, DNS will automatically be disabled. This is done by disabling the domain lookup status.

Parameters

These parameters are displayed:

Name Server IP Address – Specifies the IPv4 or IPv6 address of a domain name server to use for name-to-address resolution. Up to six IP addresses can be added to the name server list.

Web Interface

To create a list name servers:

1. Click IP Service, DNS.
2. Select Add Name Server from the Action list.
3. Enter one name server at a time.
4. Click Apply.

Figure 324: Configuring a List of Name Servers for DNS

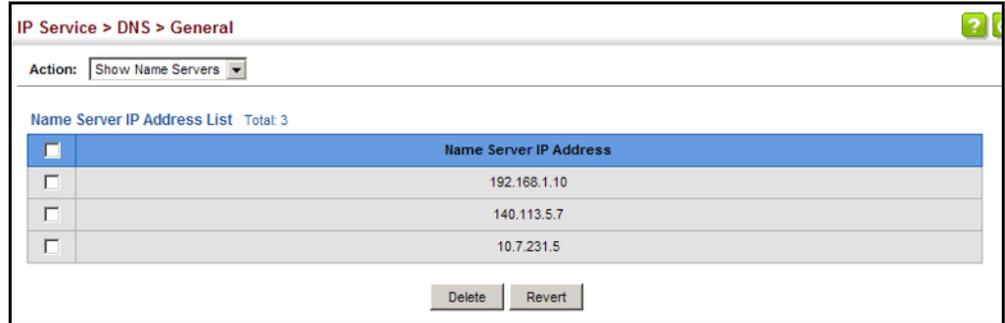


The screenshot shows a web interface for configuring DNS. At the top, the breadcrumb navigation reads "IP Service > DNS > General". Below this, there is a section for "Action:" with a dropdown menu currently set to "Add Name Server". Underneath, there is a label "Name Server IP Address" followed by a text input field containing the IP address "192.168.1.10". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

To show the list name servers:

1. Click IP Service, DNS.
2. Select Show Name Servers from the Action list.

Figure 325: Showing the List of Name Servers for DNS



Configuring Static DNS Host to Address Entries

Use the IP Service > DNS - Static Host Table (Add) page to manually configure static entries in the DNS table that are used to map domain names to IP addresses.

Command Usage

Static entries may be used for local devices connected directly to the attached network, or for commonly used resources located elsewhere on the network.

Parameters

These parameters are displayed:

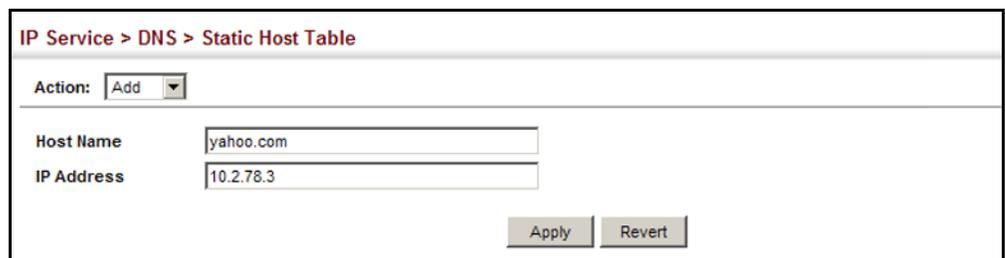
- ◆ **Host Name** – Name of a host device that is mapped to one or more IP addresses. (Range: 1-127 characters)
- ◆ **IP Address** – IPv4 or IPv6 address(es) associated with a host name.

Web Interface

To configure static entries in the DNS table:

1. Click IP Service, DNS, Static Host Table.
2. Select Add from the Action list.
3. Enter a host name and the corresponding address.
4. Click Apply.

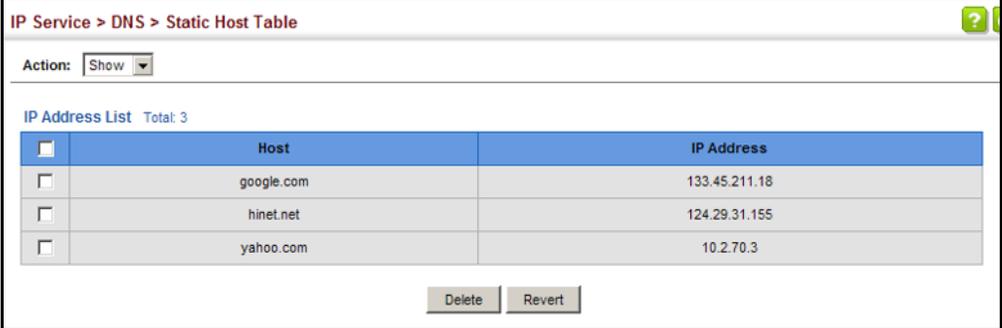
Figure 326: Configuring Static Entries in the DNS Table



To show static entries in the DNS table:

1. Click IP Service, DNS, Static Host Table.
2. Select Show from the Action list.

Figure 327: Showing Static Entries in the DNS Table



IP Service > DNS > Static Host Table

Action: Show

IP Address List Total: 3

<input type="checkbox"/>	Host	IP Address
<input type="checkbox"/>	google.com	133.45.211.18
<input type="checkbox"/>	hinet.net	124.29.31.155
<input type="checkbox"/>	yahoo.com	10.2.70.3

Delete Revert

Displaying the DNS Cache

Use the IP Service > DNS - Cache page to display entries in the DNS cache that have been learned via the designated name servers.

Command Usage

Servers or other network devices may support one or more connections via multiple IP addresses. If more than one IP address is associated with a host name via information returned from a name server, a DNS client can try each address in succession, until it establishes a connection with the target device.

Parameters

These parameters are displayed:

- ◆ **No.** – The entry number for each resource record.
- ◆ **Flag** – The flag is always “4” indicating a cache entry and therefore unreliable.
- ◆ **Type** – This field includes CNAME which specifies the host address for the owner, and ALIAS which specifies an alias.
- ◆ **IP** – The IP address associated with this record.
- ◆ **TTL** – The time to live reported by the name server.
- ◆ **Host** – The host name associated with this record.

Web Interface

To display entries in the DNS cache:

1. Click IP Service, DNS, Cache.

Figure 328: Showing Entries in the DNS Cache

The screenshot shows a web interface titled "IP Service > DNS > Cache". Below the title, it says "Cache Information Total: 3". There is a table with the following data:

No.	Flag	Type	IP	TTL	Host
1	4	CNAME	192.168.110.2	360	www.sina.com.cn
2	4	CNAME	10.2.44.3	892	www.yahoo.akadns.new
3	4	ALIAS	pointer to: 2	298	www.yahoo.com

Below the table is a "Clear" button.

Dynamic Host Configuration Protocol

Dynamic Host Configuration Protocol (DHCP) can dynamically allocate an IP address and other configuration information to network clients when they boot up. If a subnet does not already include a BOOTP or DHCP server, you can relay DHCP client requests to a DHCP server on another subnet.

Specifying A DHCP Client Identifier

Use the IP Service > DHCP > Client page to specify the DHCP client identifier for a VLAN interface.

Command Usage

- ◆ The class identifier is used identify the vendor class and configuration of the switch to the DHCP server, which then uses this information to decide on how to service the client or the type of information to return.
- ◆ The general framework for this DHCP option is set out in RFC 2132 (Option 60). This information is used to convey configuration settings or other identification information about a client, but the specific string to use should be supplied by your service provider or network administrator. Options 60, 66 and 67 statements can be added to the server daemon's configuration file.

Table 39: Options 60, 66 and 67 Statements

Option	Statement	
	Keyword	Parameter
60	vendor-class-identifier	a string indicating the vendor class identifier
66	tftp-server-name	a string indicating the tftp server name
67	bootfile-name	a string indicating the bootfile name

- ◆ By default, DHCP option 66/67 parameters are not carried in a DHCP server reply. To ask for a DHCP reply with option 66/67 information, the DHCP client request sent by this switch includes a “parameter request list” asking for this information. Besides, the client request also includes a “vendor class identifier” that allows the DHCP server to identify the device, and select the appropriate configuration file for download. This information is included in Option 55 and 124.

Table 40: Options 55 and 124 Statements

Option	Statement	
	Keyword	Parameter
55	dhcp-parameter-request-list	a list of parameters, separated by "
124	vendor-class-identifier	a string indicating the vendor class identifier

- ◆ The server should reply with the TFTP server name and boot file name.
- ◆ Note that the vendor class identifier can be formatted in either text or hexadecimal, but the format used by both the client and server must be the same.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of configured VLAN.
- ◆ **Vendor Class ID** – The following options are supported when the check box is marked to enable this feature:
 - **Default** – The default string is the model number.
 - **Text** – A text string. (Range: 1-32 characters)
 - **Hex** – A hexadecimal value.

Web Interface

To configure a DHCP client identifier:

1. Click IP Service, DHCP, Client.
2. Mark the check box to enable this feature. Select the default setting, or the format for a vendor class identifier. If a non-default value is used, enter a text string or hexadecimal value.
3. Click Apply.

Figure 329: Specifying a DHCP Client Identifier

IP Service > DHCP > Client

VLAN: 1

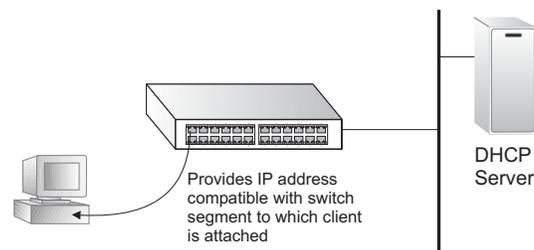
Vendor Class ID: Default AOS5700-54X

Apply Revert

Configuring DHCP Relay Service

Use the IP Service > DHCP > Relay page to configure DHCP relay service for attached host devices. If DHCP relay is enabled, and this switch sees a DHCP request broadcast, it inserts its own IP address into the request so that the DHCP server will know the subnet where the client is located. Then, the switch forwards the packet to the DHCP server. When the server receives the DHCP request, it allocates a free IP address for the DHCP client from its defined scope for the DHCP client's subnet, and sends a DHCP response back to the DHCP relay agent (i.e., this switch). This switch then passes the DHCP response received from the server to the client.

Figure 330: Layer 3 DHCP Relay Service



Command Usage

- ◆ You must specify the IP address for at least one active DHCP server. Otherwise, the switch's DHCP relay agent will not be able to forward client requests to a DHCP server. Up to five DHCP servers can be specified in order of preference.
- ◆ DHCP relay configuration will be disabled if an active DHCP server is detected on the same network segment.

Parameters

These parameters are displayed:

- ◆ **VLAN ID** – ID of configured VLAN.
- ◆ **Server IP Address** – Addresses of DHCP servers or relay servers to be used by the switch's DHCP relay agent in order of preference.
- ◆ **Restart DHCP Relay** – Use this button to re-initialize DHCP relay service.

Web Interface

To configure DHCP relay service:

1. Click IP Service, DHCP, Relay.
2. Enter up to five IP addresses for any VLAN.
3. Click Apply.

Figure 331: Configuring DHCP Relay Service

IP Service > DHCP > Relay

Note: DHCP relay configuration will be disabled if an active DHCP server is detected on the same network segment.

DHCP Server by VLAN List Total: 1

VLAN	Server IP Address				
1	<input type="text" value="192.168.2.33"/>	<input type="text" value="0.0.0.0"/>	<input type="text" value="0.0.0.0"/>	<input type="text" value="0.0.0.0"/>	<input type="text" value="0.0.0.0"/>

Click the button to restart DHCP Relay service.

General IP Routing

This chapter provides information on network functions including:

- ◆ [Ping](#) – Sends ping message to another node on the network.
- ◆ [Trace](#) – Sends ICMP echo request packets to another node on the network.
- ◆ [Address Resolution Protocol](#) – Describes how to configure ARP aging time, proxy ARP, or static addresses. Also shows how to display dynamic entries in the ARP cache.
- ◆ [Static Routes](#) – Configures static routes to other network segments.
- ◆ [Routing Table](#) – Displays routing entries learned through dynamic routing and statically configured entries.
- ◆ [Equal-cost Multipath Routing](#) – Configures the maximum number of equal-cost paths that can transmit traffic to the same destination

Overview

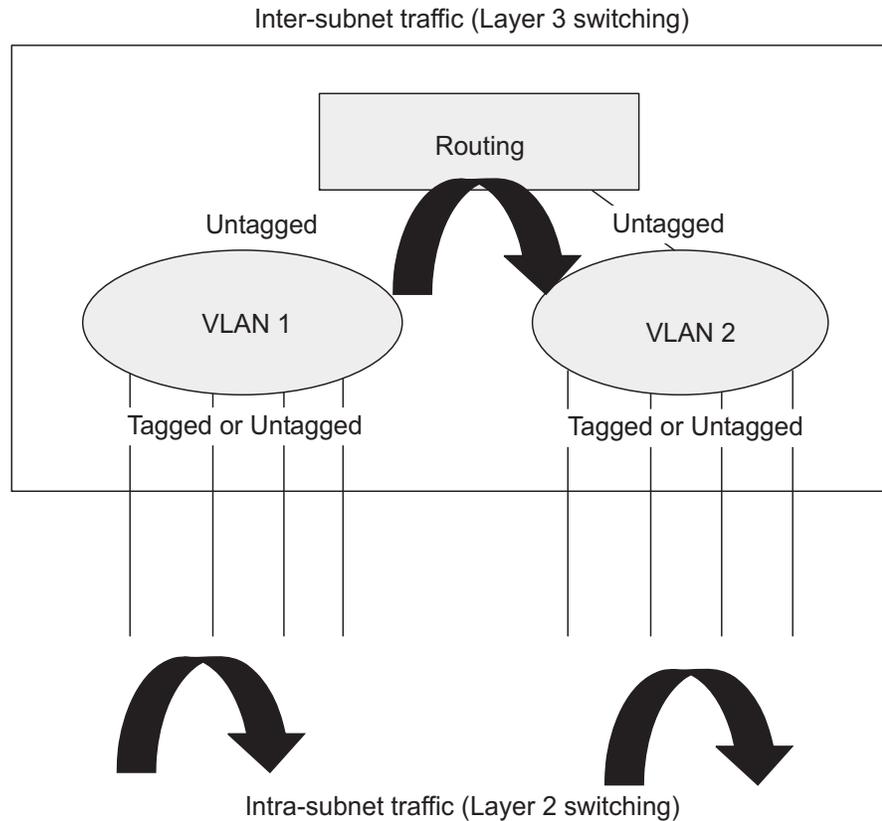
This switch supports IP routing and routing path management via static routing definitions and dynamic routing protocols such as RIP, OSPFv2, OSPFv3¹⁶, or BGPv4¹⁶. When IP routing is functioning, this switch acts as a wire-speed router, passing traffic between VLANs with different IP interfaces, and routing traffic to external IP networks. However, when the switch is first booted, default routing can only forward traffic between local IP interfaces. As with all traditional routers, static and dynamic routing functions must first be configured to work.

Initial Configuration By default, all ports belong to the same VLAN and the switch provides only Layer 2 functionality. To segment the attached network, first create VLANs for each unique user group or application traffic ([page 147](#)), assign all ports that belong to the same group to these VLANs ([page 150](#)), and then assign an IP interface to each VLAN ([page 518](#)). By separating the network into different VLANs, it can be partitioned into subnetworks that are disconnected at Layer 2. Network traffic within the same subnet is still switched using Layer 2 switching. And the VLANs can now be interconnected (as required) with Layer 3 switching.

16. Refer to the *Command Reference Guide* for information on configuring these protocols.

Each VLAN represents a virtual interface to Layer 3. You just need to provide the network address for each virtual interface, and the traffic between different subnetworks will be routed by Layer 3 switching.

Figure 332: Virtual Interfaces and Layer 3 Routing



IP Routing and Switching

IP Switching (or packet forwarding) encompasses tasks required to forward packets for both Layer 2 and Layer 3, as well as traditional routing. These functions include:

- ◆ Layer 2 forwarding (switching) based on the Layer 2 destination MAC address
- ◆ Layer 3 forwarding (routing):
 - Based on the Layer 3 destination address
 - Replacing destination/source MAC addresses for each hop
 - Incrementing the hop count
 - Decrementing the time-to-live
 - Verifying and recalculating the Layer 3 checksum

If the destination node is on the same subnetwork as the source network, then the packet can be transmitted directly without the help of a router. However, if the MAC

address is not yet known to the switch, an Address Resolution Protocol (ARP) packet with the destination IP address is broadcast to get the destination MAC address from the destination node. The IP packet can then be sent directly with the destination MAC address.

If the destination belongs to a different subnet on this switch, the packet can be routed directly to the destination node. However, if the packet belongs to a subnet not included on this switch, then the packet should be sent to the next hop router (with the MAC address of the router itself used as the destination MAC address, and the destination IP address of the destination node). The router will then forward the packet to the destination node through the correct path. The router can also use the ARP protocol to find out the MAC address of the destination node of the next router as necessary.



Note: In order to perform IP switching, the switch should be recognized by other network nodes as an IP router, either by setting it as the default gateway or by redirection from another router via the ICMP process.

When the switch receives an IP packet addressed to its own MAC address, the packet follows the Layer 3 routing process. The destination IP address is checked against the Layer 3 address table. If the address is not already there, the switch broadcasts an ARP packet to all the ports on the destination VLAN to find out the destination MAC address. After the MAC address is discovered, the packet is reformatted and sent out to the destination. The reformat process includes decreasing the Time-To-Live (TTL) field of the IP header, recalculating the IP header checksum, and replacing the destination MAC address with either the MAC address of the destination node or that of the next hop router.

When another packet destined to the same node arrives, the destination MAC can be retrieved directly from the Layer 3 address table; the packet is then reformatted and sent out the destination port. IP switching can be done at wire-speed when the destination address entry is already in the Layer 3 address table.

If the switch determines that a frame must be routed, the route is calculated only during setup. Once the route has been determined, all packets in the current flow are simply switched or forwarded across the chosen path. This takes advantage of the high throughput and low latency of switching by enabling the traffic to bypass the routing engine once the path calculation has been performed.

Routing Path Management

Routing Path Management involves the determination and updating of all the routing information required for packet forwarding, including:

- ◆ Handling routing protocols
- ◆ Updating the routing table
- ◆ Updating the Layer 3 switching database

Routing Protocols The switch supports both static and dynamic routing.

- ◆ Static routing requires routing information to be stored in the switch either manually or when a connection is set up by an application outside the switch.
- ◆ Dynamic routing uses a routing protocol to exchange routing information, calculate routing tables, and respond to changes in the status or loading of the network.

Configuring IP Routing Interfaces

Configuring Local and Remote Interfaces Use the IP > General > Routing Interface (Add Address) page to configure routing interfaces for directly connected IPv4 subnets (see [“Setting the Switch’s IP Address \(IP Version 4\)” on page 481](#)). Or use the IP > IPv6 Configuration pages to configure routing interfaces for directly connected IPv6 subnets (see [“Setting the Switch’s IP Address \(IP Version 6\)” on page 485](#)).

If this router is directly connected to end node devices (or connected to end nodes through shared media) that will be assigned to a specific subnet, then you must create a router interface for each VLAN that will support routing. The router interface consists of an IP address and subnet mask. This interface address defines both the network prefix number to which the router interface is attached and the router’s host number on that network. In other words, a router interface address defines the network segment that is connected to that interface, and allows you to send IP packets to or from the router.

You can specify the IP subnets connected directly to this router by manually assigning an IP address to each VLAN or using BOOTP or DHCP to dynamically assign an address. To specify IP subnets not directly connected to this router, you can either configure static routes (see [page 526](#)), or use the RIP, OSPF or OSPFv3 dynamic routing protocols (see [page 543](#)) to identify routes that lead to other interfaces by exchanging protocol messages with other routers on the network.

Once IP interfaces have been configured, the switch functions as a multilayer routing switch, operating at either Layer 2 or 3 as required. All IP packets are routed directly between local interfaces, or indirectly to remote interfaces using either static or dynamic routing. All other packets for non-IP protocols (for example, NetBuei, NetWare or AppleTalk) are switched based on MAC addresses).

To route traffic between remote IP interfaces, the switch should be recognized by other network nodes as an IP router, either by setting it to advertise itself as the default gateway or by redirection from another router via the ICMP process used by various routing protocols.

If the switch is configured to advertise itself as the default gateway, a routing protocol must still be used to determine the next hop router for any unknown

destinations, i.e., packets that do not match any routing table entry. If another router is designated as the default gateway, then the switch will pass packets to this router for any unknown hosts or subnets.

To configure a default gateway for IPv4, use the static routing table as described on [page 526](#), enter 0.0.0.0 for the IP address and subnet mask, and then specify this switch itself or another router as the gateway. To configure a gateway for IPv6, see [“Configuring the IPv6 Default Gateway” on page 485](#).

Using the Ping Function Use the IP > General > Ping page to send ICMP echo request packets to another node on the network.

Parameters

These parameters are displayed:

- ◆ **Host Name/IP Address** – IPv4/IPv6 address or alias of the host.
For host name-to-IP address translation to function properly, host name lookup must be enabled ([“Configuring General DNS Service Parameters” on page 505](#)), and one or more DNS servers specified (see [“Configuring a List of Name Servers” on page 508](#), or [“Configuring Static DNS Host to Address Entries” on page 509](#)).
- ◆ **Probe Count** – Number of packets to send. (Range: 1-16)
- ◆ **Packet Size** – Number of bytes in a packet. (Range: 32-512 bytes for IPv4, 0-1500 bytes for IPv6)
The actual packet size will be eight bytes larger than the size specified because the switch adds header information.

Command Usage

- ◆ Use the ping command to see if another site on the network can be reached.
- ◆ The following are some results of the **ping** command:
 - *Normal response* - The normal response occurs in one to ten seconds, depending on network traffic.
 - *Destination does not respond* - If the host does not respond, a “timeout” appears in ten seconds.
 - *Destination unreachable* - The gateway for this destination indicates that the destination is unreachable.
 - *Network or host unreachable* - The gateway found no corresponding entry in the route table.
- ◆ The same link-local address may be used by different interfaces/nodes in different zones (RFC 4007). Therefore, when specifying a link-local address,

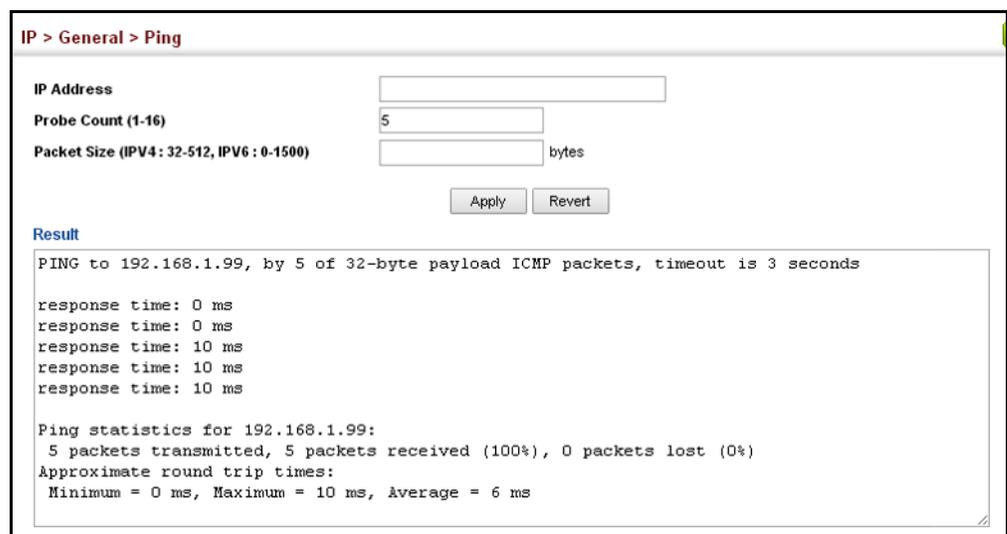
include zone-id information indicating the VLAN identifier after the % delimiter. For example, FE80::7272%1 identifies VLAN 1 as the interface.

Web Interface

To ping another device on the network:

1. Click IP, General, Ping.
2. Specify the target device and ping parameters.
3. Click Apply.

Figure 333: Pinging a Network Device



Using the Trace Route Function Use the IP > General > Trace Route page to show the route packets take to the specified destination.

Parameters

These parameters are displayed:

- ◆ **Destination IP Address** – IPv4/IPv6 address of the host.
- ◆ **IPv4 Max Failures** – The maximum number of failures before which the trace route is terminated. (Fixed: 5)
- ◆ **IPv6 Max Failures** – The maximum number of failures before which the trace route is terminated. (Range: 1-255; Default: 5)

Command Usage

- ◆ Use the trace route function to determine the path taken to reach a specified destination.

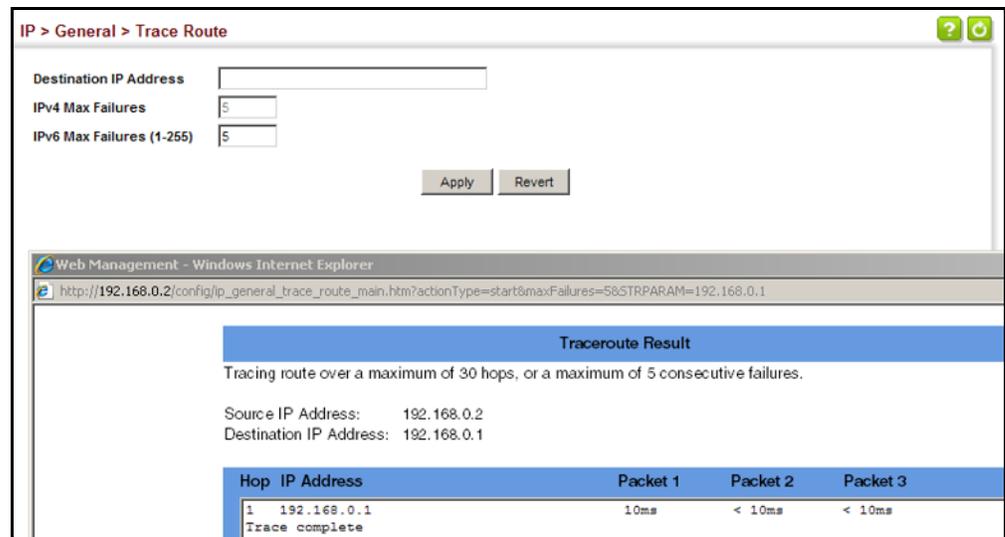
- ◆ A trace terminates when the destination responds, when the maximum timeout (TTL) is exceeded, or the maximum number of hops is exceeded.
- ◆ The trace route function first sends probe datagrams with the TTL value set at one. This causes the first router to discard the datagram and return an error message. The trace function then sends several probe messages at each subsequent TTL level and displays the round-trip time for each message. Not all devices respond correctly to probes by returning an “ICMP port unreachable” message. If the timer goes off before a response is returned, the trace function prints a series of asterisks and the “Request Timed Out” message. A long sequence of these messages, terminating only when the maximum timeout has been reached, may indicate this problem with the target device.
- ◆ The same link-local address may be used by different interfaces/nodes in different zones (RFC 4007). Therefore, when specifying a link-local address, include zone-id information indicating the VLAN identifier after the % delimiter. For example, FE80::7272%1 identifies VLAN 1 as the interface from which the trace route is sent. Note that the zone-id for the craft interface is 4097.

Web Interface

To trace the route to another device on the network:

1. Click IP, General, Trace Route.
2. Specify the target device.
3. Click Apply.

Figure 334: Tracing the Route to a Network Device



Address Resolution Protocol

If IP routing is enabled (page 543), the router uses its routing tables to make routing decisions, and uses Address Resolution Protocol (ARP) to forward traffic from one hop to the next. ARP is used to map an IP address to a physical layer (i.e., MAC) address. When an IP frame is received by this router (or any standards-based router), it first looks up the MAC address corresponding to the destination IP address in the ARP cache. If the address is found, the router writes the MAC address into the appropriate field in the frame header, and forwards the frame on to the next hop. IP traffic passes along the path to its final destination in this way, with each routing device mapping the destination IP address to the MAC address of the next hop toward the recipient, until the packet is delivered to the final destination.

If there is no entry for an IP address in the ARP cache, the router will broadcast an ARP request packet to all devices on the network. The ARP request contains the following fields similar to that shown in this example:

Table 41: Address Resolution Protocol

destination IP address	10.1.0.19
destination MAC address	?
source IP address	10.1.0.253
source MAC address	00-00-ab-cd-00-00

When devices receive this request, they discard it if their address does not match the destination IP address in the message. However, if it does match, they write their own hardware address into the destination MAC address field and send the message back to the source hardware address. When the source device receives a reply, it writes the destination IP address and corresponding MAC address into its cache, and forwards the IP traffic on to the next hop. As long as this entry has not timed out, the router will be able forward traffic directly to the next hop for this destination without having to broadcast another ARP request.

Also, if the switch receives a request for its own IP address, it will send back a response, and also cache the MAC of the source device's IP address.

ARP Timeout Configuration Use the IP > ARP (Configure General) page to specify the timeout for ARP cache entries.

Parameters

These parameters are displayed:

- ◆ **Timeout** – Sets the aging time for dynamic entries in the ARP cache. (Range: 300 - 86400 seconds; Default: 1200 seconds or 20 minutes)

The ARP aging timeout can be set for any configured VLAN.

The aging time determines how long dynamic entries remain in the cache. If the timeout is too short, the router may tie up resources by repeating ARP requests for addresses recently flushed from the table.

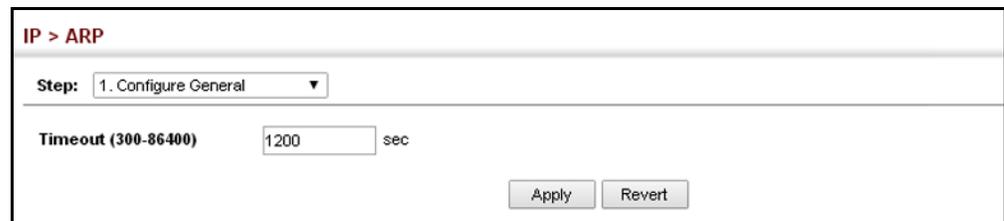
When a ARP entry expires, it is deleted from the cache and an ARP request packet is sent to re-establish the MAC address.

Web Interface

To configure the timeout for the ARP cache:

1. Click IP, ARP.
2. Select Configure General from the Step List.
3. Set the timeout to a suitable value for the ARP cache.
4. Click Apply.

Figure 335: Configuring ARP Timeout



The screenshot shows the configuration page for IP > ARP. The page title is "IP > ARP". Below the title, there is a "Step:" dropdown menu set to "1. Configure General". The main configuration area has a label "Timeout (300-86400)" followed by a text input field containing the value "1200" and the unit "sec". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

Configuring Static ARP Addresses

For devices that do not respond to ARP requests or do not respond in a timely manner, traffic will be dropped because the IP address cannot be mapped to a physical address. If this occurs, use the IP > ARP (Configure Static Address – Add) page to manually map an IP address to the corresponding physical address in the ARP cache.

Command Usage

- ◆ The ARP cache is used to map 32-bit IP addresses into 48-bit hardware (that is, Media Access Control) addresses. This cache includes entries for hosts and other routers on local network interfaces defined on this router.
- ◆ You can define up to 128 static entries in the ARP cache.
- ◆ A static entry may need to be used if there is no response to an ARP broadcast message. For example, some applications may not respond to ARP requests or the response arrives too late, causing network operations to time out.
- ◆ Static entries will not be aged out or deleted when power is reset. You can only remove a static entry via the configuration interface.

- ◆ Static entries are only displayed on the Show page for VLANs that are up. In other words, static entries are only displayed when configured for the IP subnet of a existing VLAN, and that VLAN is linked up.

Parameters

These parameters are displayed:

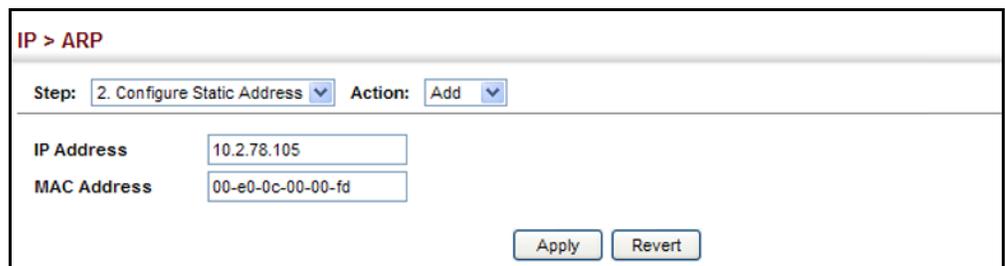
- ◆ **IP Address** – IP address statically mapped to a physical MAC address. (Valid IP addresses consist of four numbers, 0 to 255, separated by periods.)
- ◆ **MAC Address** – MAC address statically mapped to the corresponding IP address. (Valid MAC addresses are hexadecimal numbers in the format: xx-xx-xx-xx-xx-xx)

Web Interface

To map an IP address to the corresponding physical address in the ARP cache:

1. Click IP, ARP.
2. Select Configure Static Address from the Step List.
3. Select Add from the Action List.
4. Enter the IP address and the corresponding MAC address.
5. Click Apply.

Figure 336: Configuring Static ARP Entries

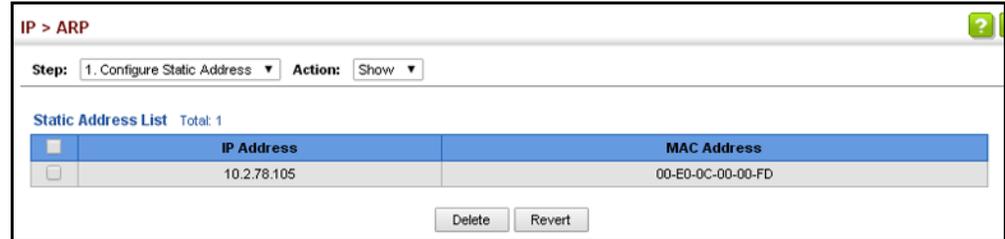


The screenshot shows a web interface for configuring static ARP entries. The breadcrumb path is "IP > ARP". The "Step" dropdown is set to "2. Configure Static Address" and the "Action" dropdown is set to "Add". There are two input fields: "IP Address" with the value "10.2.78.105" and "MAC Address" with the value "00-e0-0c-00-00-fd". At the bottom right, there are "Apply" and "Revert" buttons.

To display static entries in the ARP cache:

1. Click IP, ARP.
2. Select Configure Static Address from the Step List.
3. Select Show from the Action List.

Figure 337: Displaying Static ARP Entries



Displaying Dynamic or Local ARP Entries

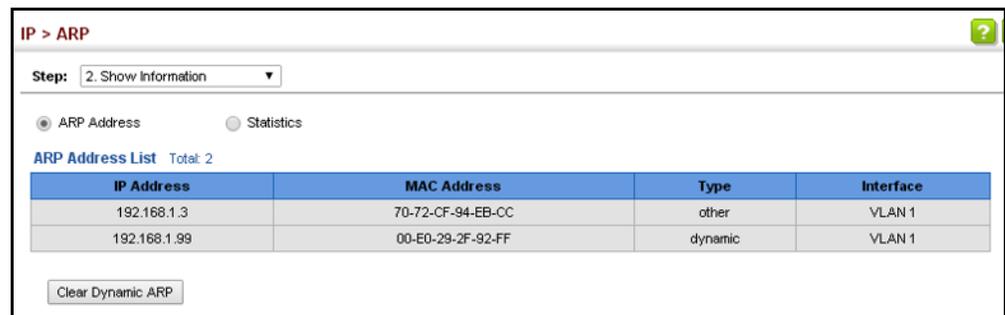
The ARP cache contains static entries, and entries for local interfaces, including subnet, host, and broadcast addresses. However, most entries will be dynamically learned through replies to broadcast messages. Use the IP > ARP (Show Information) page to display dynamic or local entries in the ARP cache.

Web Interface

To display all entries in the ARP cache:

1. Click IP, ARP.
2. Select Show Information from the Step List.
3. Click ARP Address.

Figure 338: Displaying ARP Entries



Displaying ARP Statistics

Use the IP > ARP (Show Information) page to display statistics for ARP messages crossing all interfaces on this router.

Parameters

These parameters are displayed:

Table 42: ARP Statistics

Parameter	Description
Received Request	Number of ARP Request packets received by the router.
Received Reply	Number of ARP Reply packets received by the router.

Table 42: ARP Statistics (Continued)

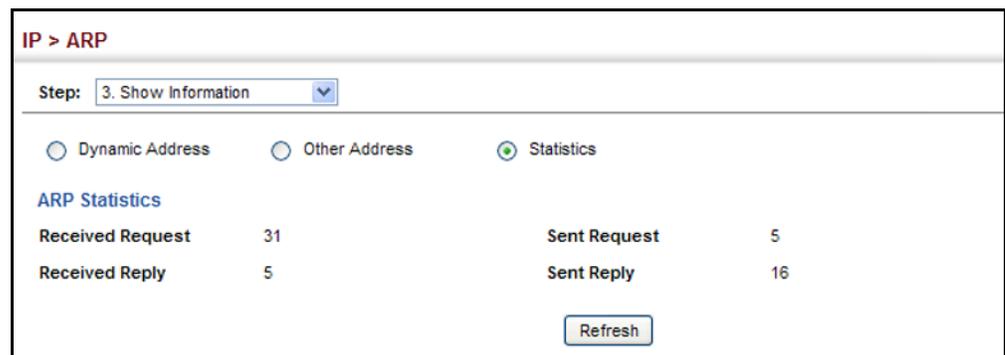
Parameter	Description
Sent Request	Number of ARP Request packets sent by the router.
Sent Reply	Number of ARP Reply packets sent by the router.

Web Interface

To display ARP statistics:

1. Click IP, ARP.
2. Select Show Information from the Step List.
3. Click Statistics.

Figure 339: Displaying ARP Statistics



Configuring Static Routes

This router can dynamically configure routes to other network segments using dynamic routing protocols (i.e., RIP or OSPF). However, you can also manually enter static routes in the routing table using the IP > Routing > Static Routes (Add) page. Static routes may be required to access network segments where dynamic routing is not supported, or can be set to force the use of a specific route to a subnet, rather than using dynamic routing. Static routes do not automatically change in response to changes in network topology, so you should only configure a small number of stable routes to ensure network accessibility.

Command Usage

- ◆ Up to 256 static routes can be configured.
- ◆ Up to eight equal-cost multipaths (ECMP) can be configured for static routing (see [“Equal-cost Multipath Routing” on page 529](#)).
- ◆ If an administrative distance is defined for a static route, and the same destination can be reached through a dynamic route at a lower administration distance, then the dynamic route will be used.

- ◆ If both static and dynamic paths have the same lowest cost, the first route stored in the routing table, either statically configured or dynamically learned via a routing protocol, will be used.
- ◆ Static routes are included in RIP and OSPF updates periodically sent by the router if this feature is enabled by (see page 553 or 580 respectively).

Parameters

These parameters are displayed:

- ◆ **Destination IP Address** – IP address of the destination network, subnetwork, or host.
- ◆ **Net Mask / Prefix Length** – Network mask for the associated IP subnet. This mask identifies the host address bits used for routing to specific subnets.
- ◆ **Next Hop** – IP address of the next router hop used for this route.
- ◆ **Distance** – An administrative distance indicating that this route can be overridden by dynamic routing information if the distance of the dynamic route is less than that configured for the static route. Note that the default administrative distances used by the dynamic unicast routing protocols is 110 for OSPF and 120 for RIP. (Range: 1-255, Default: 1)

Web Interface

To configure static routes:

1. Click IP, Routing, Static Routes.
2. Select Add from the Action List.
3. Enter the destination address, subnet mask, and next hop router.
4. Click Apply.

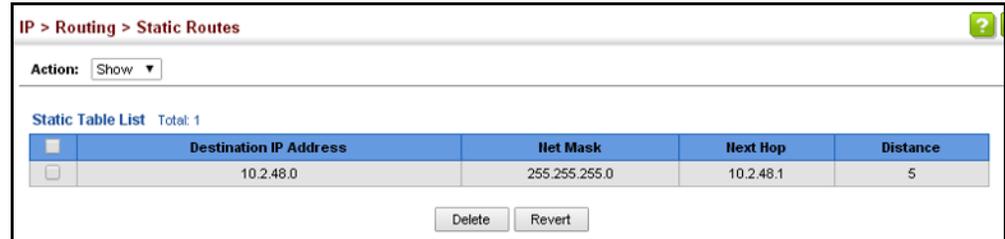
Figure 340: Configuring Static Routes

IP > Routing > Static Routes	
Action:	Add ▼
Destination IP Address	<input type="text" value="10.2.48.0"/>
Net Mask	<input type="text" value="255.255.255.0"/>
Next Hop	<input type="text" value="10.2.48.1"/>
Distance (1-255)	<input type="text" value="5"/> (Optional)
<input type="button" value="Apply"/> <input type="button" value="Revert"/>	

To display static routes:

1. Click IP, Routing, Static Routes.
2. Select Show from the Action List.

Figure 341: Displaying Static Routes



IP > Routing > Static Routes

Action: Show

Static Table List Total: 1

<input type="checkbox"/>	Destination IP Address	Net Mask	Next Hop	Distance
<input type="checkbox"/>	10.2.48.0	255.255.255.0	10.2.48.1	5

Delete Revert

Displaying the Routing Table

Use the IP > Routing > Routing Table (Show Information) page to display all routes that can be accessed via local network interfaces, through static routes, or through a dynamically learned route. If route information is available through more than one of these methods, the priority for route selection is local, static, and then dynamic (except when the distance parameter of a dynamic route is set to a value that makes its priority exceed that of a static route). Also note that the route for a local interface is not enabled (i.e., listed in the routing table) unless there is at least one active link connected to that interface.

Command Usage

- ◆ The Forwarding Information Base (FIB) contains information required to forward IP traffic. It contains the interface identifier and next hop information for each reachable destination network prefix based on the IP routing table. When routing or topology changes occur in the network, the routing table is updated, and those changes are immediately reflected in the FIB.

The FIB is distinct from the routing table (or, Routing Information Base – RIB), which holds all routing information received from routing peers. The FIB contains unique paths only. It does not contain any secondary paths. A FIB entry consists of the minimum amount of information necessary to make a forwarding decision on a particular packet. The typical components within a FIB entry are a network prefix, a router (i.e., VLAN) interface, and next hop information.
- ◆ The Routing Table only displays routes which are currently accessible for forwarding. The router must be able to directly reach the next hop, so the VLAN interface associated with any dynamic or static route entry must be up. Note that routes currently not accessible for forwarding, may still be displayed by using the “show ip route database” command described in the *CLI Reference Guide*.

Parameters

These parameters are displayed:

- ◆ **VLAN** – VLAN identifier (i.e., configure as a valid IP subnet).
- ◆ **Destination IP Address** – IP address of the destination network, subnetwork, or host. Note that the address 0.0.0.0 indicates the default gateway for this router.
- ◆ **Net Mask / Prefix Length** – Network mask for the associated IP subnet. This mask identifies the host address bits used for routing to specific subnets.
- ◆ **Next Hop** – The IP address of the next hop (or gateway) in this route.
- ◆ **Metric** – Cost for this interface.
- ◆ **Protocol** – The protocol which generated this route information. (Options: Local, Static, RIP, OSPF, Others)

Web Interface

To display the routing table:

1. Click IP, Routing, Routing Table.
2. Select Show Information from the Action List.

Figure 342: Displaying the Routing Table

The screenshot shows the 'IP > Routing > Routing Table' page. At the top, there is a breadcrumb trail and an 'Action:' dropdown menu set to 'Show Information'. Below this is a 'Routing Table List' with a 'Total: 5' indicator. The table has six columns: VLAN, Destination IP Address, Net Mask / Prefix Length, Next Hop, Metric, and Protocol. The data rows are as follows:

VLAN	Destination IP Address	Net Mask / Prefix Length	Next Hop	Metric	Protocol
0	127.0.0.0	255.0.0.0	--	0	Local
2	192.168.0.0	255.255.255.0	--	0	Local
1	192.168.2.0	255.255.255.0	--	0	Local
2	192.168.3.0	255.255.255.0	192.168.0.1	0	Static
0	:::1	128	--	0	Local

Equal-cost Multipath Routing

Use the IP > Routing > Routing Table (Configure ECMP Number) page to configure the maximum number of equal-cost paths that can transmit traffic to the same destination. The Equal-cost Multipath routing algorithm is a technique that supports load sharing over multiple equal-cost paths for data passing to the same destination. Whenever multiple paths with equal path cost to the same destination are found in the routing table, the ECMP algorithm first checks if the cost is lower than that of any other entries in the routing table. If the cost is the lowest in the table, the switch will use up to eight of the paths with equal lowest cost to balance

the traffic forwarded to the destination. ECMP uses either equal-cost multipaths manually configured in the static routing table, or equal-cost multipaths dynamically generated by the Open Shortest Path Algorithm (OSPF). In other words, it uses either static or OSPF entries, not both. Normal unicast routing simply selects the path to the destination that has the lowest cost. Multipath routing still selects the path with the lowest cost, but can forward traffic over multiple paths if they all have the same lowest cost. ECMP is enabled by default on the switch. If there is only one lowest cost path toward the destination, this path will be used to forward all traffic. If there is more than one lowest-cost path configured in the static routing table (see [“Configuring Static Routes” on page 526](#)), or dynamically generated by OSPFv2 (see [“Configuring the Open Shortest Path First Protocol \(Version 2\)” on page 562](#)), then up to 8 paths with the same lowest cost can be used to forward traffic to the destination.

Command Usage

- ◆ ECMP only selects paths of the same protocol type. It cannot be applied to both static paths and dynamic paths at the same time for the same destination. If both static and dynamic paths have the same lowest cost, the static paths have precedence over dynamic paths.
- ◆ Each path toward the same destination with equal-cost takes up one entry in the routing table to record routing information. In other words, a route with 8 paths will take up 8 entries.
- ◆ The routing table can only have up to 8 equal-cost multipaths for static routing and 8 for dynamic routing for a common destination. However, the system supports up to 256 total ECMP entries in ASIC for fast switching, with any additional entries handled by software routing.
- ◆ When there are multiple paths toward the same destination with equal-cost, the system chooses one of these paths to forward each packet toward the destination by applying a load-splitting algorithm.

A hash value is calculated based upon the source and destination IP fields of each packet as an indirect index to one of the multiple paths. Because the hash algorithm is calculated based upon the packet header information which can identify specific traffic flows, this technique minimizes the number of times a path is changed for individual flows. In general, path changes for individual flows will only occur when a path is added or removed from the multipath group.

Parameters

These parameters are displayed:

- ◆ **ECMP Number** – Sets the maximum number of equal-cost paths to the same destination that can be installed in the routing table. (Range: 1-8; Default: 8)

Web Interface

To configure the maximum ECMP number:

1. Click IP, Routing, Routing Table.
2. Select Configure ECMP Number from the Action List.
3. Enter the maximum number of equal-cost paths used to route traffic to the same destination that are permitted on the switch.
4. Click Apply

Figure 343: Setting the Maximum ECMP Number



The screenshot shows a web interface for configuring the maximum ECMP number. The breadcrumb navigation is "IP > Routing > Routing Table". The "Action:" dropdown menu is set to "Configure ECMP Number". The "ECMP Number (1-8)" input field contains the value "4". There are "Apply" and "Revert" buttons at the bottom right.

IP > Routing > Routing Table	
Action:	Configure ECMP Number
ECMP Number (1-8)	4
<input type="button" value="Apply"/> <input type="button" value="Revert"/>	

Unicast Routing

This chapter describes how to configure the following unicast routing protocols:

RIP – Configures Routing Information Protocol.

OSPFv2 – Configures Open Shortest Path First (Version 2) for IPv4.

Overview

This switch can route unicast traffic to different subnetworks using Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) protocol. It supports RIP, RIP-2 and OSPFv2 dynamic routing in the web management interface. These protocols exchange routing information, calculate routing tables, and can respond to changes in the status or loading of the network. For information on configuring OSPFv3 or BGPv4, refer to the *CLI Reference Guide*.

RIP and RIP-2 Dynamic Routing Protocols

The RIP protocol is the most widely used routing protocol. RIP uses a distance-vector-based approach to routing. Routes are determined on the basis of minimizing the distance vector, or hop count, which serves as a rough estimate of transmission cost. Each router broadcasts its advertisement every 30 seconds, together with any updates to its routing table. This allows all routers on the network to learn consistent tables of next hop links which lead to relevant subnets.

OSPFv2 Dynamic Routing Protocols

OSPF overcomes all the problems of RIP. It uses a link state routing protocol to generate a shortest-path tree, then builds up its routing table based on this tree. OSPF produces a more stable network because the participating routers act on network changes predictably and simultaneously, converging on the best route more quickly than RIP. Moreover, when several equal-cost routes to a destination exist, traffic can be distributed equally among them.

Non-IP Protocol Routing

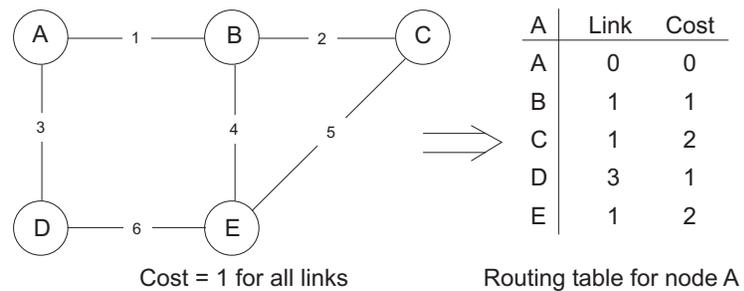
The switch supports IP routing only. Non-IP protocols such as IPX and Appletalk cannot be routed by this switch, and will be confined within their local VLAN group unless bridged by an external router.

To coexist with a network built on multilayer switches, the subnetworks for non-IP protocols must follow the same logical boundary as that of the IP subnetworks. A separate multi-protocol router can then be used to link the subnetworks by connecting to one port from each available VLAN on the network.

Configuring the Routing Information Protocol

The RIP protocol is the most widely used routing protocol. The RIP protocol uses a distance-vector-based approach to routing. Routes are determined on the basis of minimizing the distance vector, or hop count, which serves as a rough estimate of transmission cost. Each router broadcasts its advertisement every 30 seconds, together with any updates to its routing table. This allows all routers on the network to learn consistent tables of next hop links which lead to relevant subnets.

Figure 354: Configuring RIP



Command Usage

- ◆ Just as Layer 2 switches use the Spanning Tree Algorithm to prevent loops, routers also use methods for preventing loops that would cause endless retransmission of data traffic. RIP utilizes the following three methods to prevent loops from occurring:
 - Split horizon – Never propagate routes back to an interface port from which they have been acquired.
 - Poison reverse – Propagate routes back to an interface port from which they have been acquired, but set the distance-vector metrics to infinity. (This provides faster convergence.)
 - Triggered updates – Whenever a route gets changed, broadcast an update message after waiting for a short random delay, but without waiting for the periodic cycle.
- ◆ RIP-2 is a compatible upgrade to RIP. RIP-2 adds useful capabilities for plain text authentication, multiple independent RIP domains, variable length subnet masks, and multicast transmissions for route advertising (RFC 1723).
- ◆ There are several serious problems with RIP that you should consider. First of all, RIP (version 1) has no knowledge of subnets, both RIP versions can take a long time to converge on a new route after the failure of a link or router during which time routing loops may occur, and its small hop count limitation of 15 restricts its use to smaller networks. Moreover, RIP (version 1) wastes valuable network bandwidth by propagating routing information via broadcasts; it also considers too few network variables to make the best routing decision.

Configuring General Protocol Settings Use the Routing Protocol > RIP > General (Configure) page to configure general settings and the basic timers.

RIP is used to specify how routers exchange routing information. When RIP is enabled on this router, it sends RIP messages to all devices in the network every 30 seconds (by default), and updates its own routing table when RIP messages are received from other routers. To communicate properly with other routers using RIP, you need to specify the RIP version used globally by the router, as well as the RIP send and receive versions used on specific interfaces (see [“Configuring Network Interfaces for RIP” on page 556](#)).

Command Usage

- ◆ RIP is used to specify how routers exchange routing information. When RIP is enabled on this router, it sends RIP messages to all devices in the network every 30 seconds (by default), and updates its own routing table when RIP messages are received from other routers. To communicate properly with other routers using RIP, you need to specify the RIP version used globally by the router, as well as the RIP send and receive versions used on specific interfaces ([page 556](#)).

Parameters

These parameters are displayed:

Global Settings

- ◆ **RIP Routing Process** – Enables RIP routing globally. RIP must also be enabled on each network interface which will participate in the routing process as described under [“Specifying Network Interfaces” on page 549](#). (Default: Disabled)
- ◆ **Global RIP Version** – Specifies a RIP version used globally by the router. (Version 1, Version 2, By Interface; Default: By Interface)

When a Global RIP Version is specified, any VLAN interface not previously set to a specific Receive or Send Version ([page 556](#)) is set to the following values:

- RIP Version 1 configures previously unset interfaces to send RIPv1 compatible protocol messages and receive either RIPv1 or RIPv2 protocol messages.
- RIP Version 2 configures previously unset interfaces to use RIPv2 for both sending and receiving protocol messages.

RIP send/receive versions set on the RIP Interface settings screen ([page 556](#)) always take precedence over the settings for the Global RIP Version. However, when the Global RIP Version is set to “By Interface,” any VLAN interface not previously set to a specific receive or send version is set to the following default values:

- Receive: Accepts RIPv1 or RIPv2 packets.
- Send: Route information is broadcast to other routers with RIPv2.

- ◆ **RIP Default Metric** – Sets the default metric assigned to external routes imported from other protocols. (Range: 1-15; Default: 1)

The default metric must be used to resolve the problem of redistributing external routes with incompatible metrics.

It is advisable to use a low metric when redistributing routes from another protocol into RIP. Using a high metric limits the usefulness of external routes redistributed into RIP. For example, if a metric of 10 is defined for redistributed routes, these routes can only be advertised to routers up to 5 hops away, at which point the metric exceeds the maximum hop count of 15. By defining a low metric of 1, traffic can follow a imported route the maximum number of hops allowed within a RIP domain. However, note that using a low metric can increase the possibility of routing loops. For example, this can occur if there are multiple redistribution points and the router learns about the same external network with a better metric from a redistribution point other than that derived from the original source.

The default metric does not override the metric value set in the Redistribute screen (see [“Configuring Route Redistribution” on page 553](#)). When a metric value has not been configured in the Redistribute screen, the default metric sets the metric value to be used for all imported external routes.

- ◆ **RIP Max Prefix** – Sets the maximum number of RIP routes which can be installed in the routing table. (Range: 1-7168; Default: 7168)
- ◆ **Default Information Originate** – Generates a default external route into the local RIP autonomous system. (Default: Disabled)

A default route is set for every Layer 3 interface where RIP is enabled. The response packet to external queries marks each active RIP interface as a default router with the IP address 0.0.0.0.

- ◆ **Default Distance** – Defines an administrative distance for external routes learned from other routing protocols. External routes are routes for which the best path is learned from a neighbor external to the local RIP autonomous system. Routes with a distance of 255 are not installed in the routing table. (Range: 1-255; Default: 120)

Administrative distance is used by the routers to select the preferred path when there are two or more different routes to the same destination from two different routing protocols. A smaller administrative distance indicates a more reliable protocol.

Use the Routing Protocol > RIP > Distance page (see [page 555](#)) to configure the distance to a specific network address, or to configure an access list that filters networks according to the IP address of the router supplying the routing information.

- ◆ **Number of Route Changes** – The number of route changes made to the IP route database by RIP.

- ◆ **Number of Queries** – The number of responses sent to RIP queries from other systems.

Basic Timer Settings



Note: The timers must be set to the same values for all routers in the network.

- ◆ **Update** – Sets the rate at which updates are sent. This is the fundamental timer used to control all basic RIP processes. (Range: 5-2147483647 seconds; Default: 30 seconds)

Setting the update timer to a short interval can cause the router to spend an excessive amount of time processing updates. On the other hand, setting it to an excessively long time will make the routing protocol less sensitive to changes in the network configuration.

- ◆ **Timeout** – Sets the time after which there have been no update messages that a route is declared dead. The route is marked inaccessible (i.e., the metric set to infinite) and advertised as unreachable. However, packets are still forwarded on this route. (Range: 90-360 seconds; Default: 180 seconds)
- ◆ **Garbage Collection** – After the *timeout* interval expires, the router waits for an interval specified by the *garbage-collection* timer before removing this entry from the routing table. This timer allows neighbors to become aware of an invalid route prior to purging. (Range: 60-240 seconds; Default: 120 seconds)

Web Interface

To configure general settings for RIP:

1. Click Routing Protocol, RIP, General.
2. Select Configure Global from the Action list.
3. Enable RIP, set the RIP version used on unset interfaces to RIPv1 or RIPv2, set the default metric assigned to external routes, set the maximum number of routes allowed by the system, and set the basic timers.
4. Click Apply.

Figure 355: Configuring General Settings for RIP

The screenshot shows the configuration page for RIP. At the top, it says "Routing Protocol > RIP > General". Below that is an "Action:" dropdown menu set to "Configure". The page is divided into two main sections: "Global" and "Basic Timer".

Global

RIP Routing Process	<input checked="" type="checkbox"/> Enabled
Global RIP Version	By interface
RIP Default Metric (1-15)	1
RIP Max Prefix (1-7168)	7168
Default Information Originate	<input type="checkbox"/> Enabled
Default Distance (1-255)	120
Number of Route Changes	0
Number of Queries	0

Basic Timer

Update (5-2147483647)	30	sec
Timeout (90-360)	180	sec
Garbage Collection (60-240)	120	sec

At the bottom right, there are "Apply" and "Revert" buttons.

Clearing Entries from the Routing Table Use the Routing Protocol > RIP > General (Clear Route) page to clear entries from the routing table based on route type or a specific network address.

Command Usage

- ◆ RIP must be enabled to activate this menu option.
- ◆ Clearing "All" types deletes all routes in the RIP table. To avoid deleting the entire RIP network, redistribute connected routes using the Routing Protocol > RIP > Redistribute screen (page 553) to make the RIP network a connected route. To delete the RIP routes learned from neighbors, but keep the RIP network intact, clear "RIP" types from the routing table.

Parameters

These parameters are displayed:

- ◆ **Clear Route By Type** – Clears entries from the RIP routing table based on the following types:
 - **All** – Deletes all entries from the routing table.
 - **Connected** – Deletes all currently connected entries.
 - **OSPF** – Deletes all entries learned through OSPF.
 - **RIP** – Deletes all entries learned through the RIP.
 - **Static** – Deletes all static entries.

- ◆ **Clear Route By Network** – Clears a specific route based on its IP address and prefix length.
 - **Network IP Address** – Deletes all related entries for the specified network address.
 - **Prefix Length** – A decimal value indicating how many contiguous bits (from the left) of the address comprise the network portion of the address.

Web Interface

To clear entries from the routing table RIP:

1. Click Routing Protocol, RIP, General.
2. Select Clear Route from the Action list.
3. When clearing routes by type, select the required type from the drop-down list. When clearing routes by network, enter a valid network address and prefix length.
4. Click Apply.

Figure 356: Clearing Entries from the Routing Table

The screenshot shows the configuration page for the Routing Protocol > RIP > General. At the top, the breadcrumb is 'Routing Protocol > RIP > General'. Below it, the 'Action:' dropdown menu is set to 'Clear Route'. Underneath, there are two radio buttons for 'Clear Route by': 'Type' (unselected) and 'Network' (selected). Below these are two input fields: 'Network IP Address' with the value '192.168.1.0' and 'Prefix Length (1-32)' with the value '24'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

Specifying Network Interfaces Use the Routing Protocol > RIP > Network (Add) page to specify the network interfaces that will be included in the RIP routing process.

Command Usage

- ◆ RIP only sends and receives updates on specified interfaces. If a network is not specified, the interfaces in that network will not be advertised in any RIP updates.
- ◆ No networks are specified by default.

Parameters

These parameters are displayed:

- ◆ **By Address** – Adds a network to the RIP routing process.
 - **Subnet Address** – IP address of a network directly connected to this router. (Default: No networks are specified)
 - **Prefix Length** – A decimal value indicating how many contiguous bits (from the left) of the address comprise the network portion of the address. This mask identifies the network address bits used for the associated routing entries.
- ◆ **By VLAN** – Adds a Layer 3 VLAN to the RIP routing process. The VLAN must be configured with an IP address. (Range: 1-4093)

Web Interface

To add a network interface to RIP:

1. Click Routing Protocol, RIP, Network.
2. Select Add from the Action list.
3. Add an interface that will participate in RIP.
4. Click Apply.

Figure 357: Adding Network Interfaces to RIP

Routing Protocol > RIP > Network

Action: Add

By IP Address VLAN

Subnet Address: 10.1.0.0

Prefix Length (1-32): 16

Apply Revert

To show the network interfaces using RIP:

1. Click Routing Protocol, RIP, Network.
2. Select Show from the Action list.
3. Click IP Address or VLAN.

Figure 358: Showing Network Interfaces Using RIP



Specifying Passive Interfaces Use the Routing Protocol > RIP > Passive Interface (Add) page to stop RIP from sending routing updates on the specified interface.

Command Usage

- ◆ Network interfaces can be configured to stop RIP broadcast and multicast messages from being sent. If the sending of routing updates is blocked on an interface, the attached subnet will still continue to be advertised to other interfaces, and updates from other routers on the specified interface will continue to be received and processed.
- ◆ This feature can be used in conjunction with the static neighbor feature (described in the next section) to control the routing updates sent to specific neighbors.

Parameters

These parameters are displayed:

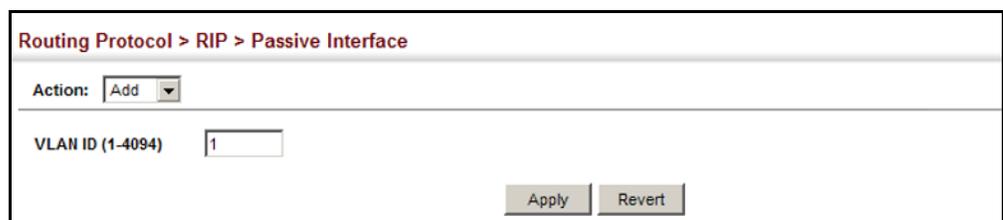
- ◆ **VLAN** – VLAN interface on which to stop sending RIP updates. (Range: 1-4094)

Web Interface

To specify a passive RIP interface:

1. Click Routing Protocol, RIP, Passive Interface.
2. Select Add from the Action list.
3. Add the interface on which to stop sending RIP updates.
4. Click Apply.

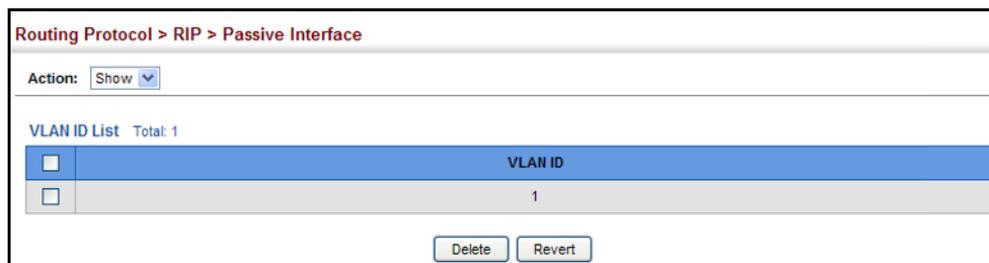
Figure 359: Specifying a Passive RIP Interface



To show the passive RIP interfaces:

1. Click Routing Protocol, RIP, Passive Interface.
2. Select Show from the Action list.

Figure 360: Showing Passive RIP Interfaces



Specifying Static Neighbors

Use the Routing Protocol > RIP > Neighbor Address (Add) page to configure this router to directly exchange routing information with a static neighbor (specifically for point-to-point links), rather than relying on broadcast or multicast messages generated by the RIP protocol. This feature can be used in conjunction with the passive interface feature (described in the preceding section) to control the routing updates sent to specific neighbors.

Parameters

These parameters are displayed:

- ◆ **IP Address** – IP address of a static neighboring router with which to exchange routing information.

Web Interface

To specify a static RIP neighbor:

1. Click Routing Protocol, RIP, Neighbor Address.
2. Select Add from the Action list.
3. Add the address of any static neighbors which may not readily be discovered through RIP.
4. Click Apply.

Figure 361: Specifying a Static RIP Neighbor

Routing Protocol > RIP > Neighbor Address

Action: Add

IP Address: 10.2.0.254

Apply Revert

To show static RIP neighbors:

1. Click Routing Protocol, RIP, Neighbor Address.
2. Select Show from the Action list.

Figure 362: Showing Static RIP Neighbors

Routing Protocol > RIP > Neighbor Address

Action: Show

Neighbor Address List Total: 1

	IP Address
<input type="checkbox"/>	10.2.0.254

Delete Revert

Configuring Route Redistribution

Use the Routing Protocol > RIP > Redistribute (Add) page to import external routing information from other routing domains (that is, directly connected routes, protocols, or static routes) into this autonomous system.

Parameters

These parameters are displayed:

- ◆ **Protocol** – The type of routes that can be imported include:
 - **BGP** - External routes will be imported from the Border Gateway Protocol (BGP) into this routing domain.
 - **Connected** – Imports routes that are established automatically just by enabling IP on an interface.
 - **Static** – Static routes will be imported into this routing domain.
 - **OSPF** – External routes will be imported from the Open Shortest Path First protocol into this routing domain.

- ◆ **Metric** – Metric assigned to all external routes for the specified protocol. (Range: 0-16; Default: the default metric as described under “Configuring General Protocol Settings” on page 545.)

A route metric must be used to resolve the problem of redistributing external routes with incompatible metrics.

When a metric value has not been configured on this page, the default-metric determines the metric value to be used for all imported external routes.

It is advisable to use a low metric when redistributing routes from another protocol into RIP. Using a high metric limits the usefulness of external routes redistributed into RIP. For example, if a metric of 10 is defined for redistributed routes, these routes can only be advertised to routers up to 5 hops away, at which point the metric exceeds the maximum hop count of 15. By defining a low metric of 1, traffic can follow an imported route the maximum number of hops allowed within a RIP domain. However, using a low metric can increase the possibility of routing loops. For example, this can occur if there are multiple redistribution points and the router learns about the same external network with a better metric from a redistribution point other than that derived from the original source.

Web Interface

To import external routing information from other routing domains:

1. Click Routing Protocol, RIP, Redistribute.
2. Select Add from the Action list.
3. Specify the protocol types (directly connected, BGP, OSPF, or static) from which to import external routes, and the metric to assign to these routes.
4. Click Apply.

Figure 363: Redistributing External Routes into RIP



The screenshot shows a web interface for configuring RIP redistribution. The breadcrumb path is "Routing Protocol > RIP > Redistribute". Below the breadcrumb, there is an "Action:" label followed by a dropdown menu currently showing "Add". Underneath, there are two rows of configuration fields: "Protocol" with a dropdown menu showing "OSPF", and "Metric (1-16)" with a text input field containing the number "3" and the text "(Optional)" to its right. At the bottom right of the form area, there are two buttons: "Apply" and "Revert".

To show external routes imported into RIP:

1. Click Routing Protocol, RIP, Redistribute.
2. Select Show from the Action list.

Figure 364: Showing External Routes Redistributed into RIP

Routing Protocol > RIP > Redistribute

Action: Show

Redistribute List Total: 1

<input type="checkbox"/>	Protocol	Metric
<input type="checkbox"/>	OSPF	3

Delete Revert

Specifying an Administrative Distance

Use the Routing Protocol > RIP > Distance (Add) page to define an administrative distance for external routes learned from other routing protocols.

Command Usage

- ◆ Administrative distance is used by the routers to select the preferred path when there are two or more different routes to the same destination from two different routing protocols. A smaller administrative distance indicates a more reliable protocol.
- ◆ The administrative distance is applied to all routes learned for the specified network.

Parameters

These parameters are displayed:

- ◆ **Distance** – Administrative distance for external routes. External routes are routes for which the best path is learned from a neighbor external to the local RIP autonomous system. Routes with a distance of 255 are not installed in the routing table. (Range: 1-255)
- ◆ **IP Address** – IP address of a route entry.
- ◆ **Subnet Mask** – This mask identifies the host address bits used for associated routing entries.

Web Interface

To define an administrative distance for external routes learned from other routing protocols:

1. Click Routing Protocol, RIP, Distance.
2. Select Add from the Action list.
3. Enter the distance and the external route.
4. Click Apply.

Figure 365: Setting the Distance Assigned to External Routes

Routing Protocol > RIP > Distance

Action: Add

Distance (1-255) 120

IP Address 192.168.3.0

Subnet Mask 255.255.255.0

Apply Revert

To show the distance assigned to external routes learned from other routing protocols:

1. Click Routing Protocol, RIP, Distance.
2. Select Show from the Action list.

Figure 366: Showing the Distance Assigned to External Routes

Routing Protocol > RIP > Distance

Action: Show

RIP Distance List Total: 1

<input type="checkbox"/>	Distance	IP Address	Subnet Mask
<input type="checkbox"/>	120	192.168.3.0	255.255.255.0

Delete Revert

Configuring Network Interfaces for RIP

Use the Routing Protocol > RIP > Interface (Add) page to configure the send/receive version, authentication settings, and the loopback prevention method for each interface that participates in the RIP routing process.

Command Usage

Specifying Receive and Send Protocol Types

- ◆ Specify the protocol message type accepted (that is, RIP version) and the message type sent (that is, RIP version or compatibility mode) for each RIP interface.
- ◆ Setting the RIP Receive Version or Send Version for an interface overrides the global setting specified in the RIP General Settings screen (see [“Configuring General Protocol Settings”](#) on page 545).
- ◆ The Send Version can be specified based on these options:
 - Use “RIPv1” or “RIPv2” if all routers in the local network are based on RIPv1 or RIPv2, respectively.
 - Use “RIPv1 Compatible” to propagate route information by broadcasting to other routers on the network using the RIPv2 advertisement list, instead of

multicasting as normally required by RIPv2. (Using this mode allows older RIPv2 routers which only receive RIP broadcast messages to receive all of the information provided by RIPv2, including subnet mask, next hop and authentication information. (This is the default setting.)

- Use “Do Not Send” to passively monitor route information advertised by other routers attached to the network.
- ◆ The Receive Version can be specified based on these options:
 - Use “RIPv1” or “RIPv2” if all routers in the local network are based on RIPv1 or RIPv2, respectively.
 - Use “RIPv1 and RIPv2” if some routers in the local network are using RIPv2, but there are still some older routers using RIPv1. (This is the default setting.)
 - Use “Do Not Receive” if dynamic entries are not required to be added to the routing table for an interface. (For example, when only static routes are to be allowed for a specific interface.)

Protocol Message Authentication

RIPv1 is not a secure protocol. Any device sending protocol messages from UDP port 520 will be considered a router by its neighbors. Malicious or unwanted protocol messages can be easily propagated throughout the network if no authentication is required.

RIPv2 supports authentication using a simple password or MD5 key encryption. When a router is configured to exchange authentication messages, it will insert the password into all transmitted protocol packets, and check all received packets to ensure that they contain the authorized password. If any incoming protocol messages do not contain the correct password, they are simply dropped.

For authentication to function properly, both the sending and receiving interface must be configured with the same password or authentication key.

Loopback Prevention

Just as Layer 2 switches use the Spanning Tree Algorithm to prevent loops, routers also use methods for preventing loops that would cause endless retransmission of data traffic. When protocol packets are caught in a loop, links will be congested, and protocol packets may be lost. However, the network will slowly converge to the new state. RIP supports several methods which can provide faster convergence when the network topology changes and prevent most loops from occurring.

Parameters

These parameters are displayed:

- ◆ **VLAN ID** – Layer 3 VLAN interface. This interface must be configured with an IP address and have an active link. (Range: 1-4094)

- ◆ **Send Version** – The RIP version to send on an interface.
 - **RIPv1:** Sends only RIPv1 packets.
 - **RIPv2:** Sends only RIPv2 packets.
 - **RIPv1 Compatible:** Route information is broadcast to other routers with RIPv2.
 - **Do Not Send:** Does not transmit RIP updates. Passively monitors route information advertised by other routers attached to the network.

The default depends on the setting for the Global RIP Version. (See [“Configuring General Protocol Settings” on page 545.](#))

- ◆ **Receive Version** – The RIP version to receive on an interface.
 - **RIPv1:** Accepts only RIPv1 packets.
 - **RIPv2:** Accepts only RIPv2 packets.
 - **RIPv1 and RIPv2:** Accepts RIPv1 and RIPv2 packets.
 - **Do Not Receive:** Does not accept incoming RIP packets. This option does not add any dynamic entries to the routing table for an interface.

The default depends on the setting for the Global RIP Version. (See [“Configuring General Protocol Settings” on page 545.](#))

- ◆ **Authentication Type** – Specifies the type of authentication required for exchanging RIPv2 protocol messages. (Default: No Authentication)
 - **No Authentication:** No authentication is required.
 - **Simple Password:** Requires the interface to exchange routing information with other routers based on an authorized password. (Note that authentication only applies to RIPv2.)
 - **MD5:** Message Digest 5 (MD5) authentication.

MD5 is a one-way hash algorithm that takes the authentication key and produces a 128 bit message digest or “fingerprint.” This makes it computationally infeasible to produce two messages having the same message digest, or to produce any message having a given pre-specified target message digest.
- ◆ **Authentication Key** – Specifies the key to use for authenticating RIPv2 packets. For authentication to function properly, both the sending and receiving interface must use the same password. (Range: 1-16 characters, case sensitive)

- ◆ **Instability Prevention** – Specifies the method used to reduce the convergence time when the network topology changes, and to prevent RIP protocol messages from looping back to the source router.
 - **Split Horizon** – This method never propagate routes back to an interface from which they have been acquired.
 - **Poison Reverse** – This method propagates routes back to an interface from which they have been acquired, but sets the distance-vector metrics to infinity. This provides faster convergence. (This is the default setting.)
 - **None** – No loopback prevention method is employed. If a loop occurs without using any prevention method, the hop count for a route may be gradually incremented to infinity (that is, 16) before the route is deemed unreachable.

Web Interface

To network interface settings for RIP:

1. Click Routing Protocol, RIP, Interface.
2. Select Add from the Action list.
3. Select a Layer 3 VLAN interface to participate in RIP. Select the RIP protocol message types that will be received and sent. Select the RIP authentication method and password. And then set the loopback prevention method.
4. Click Apply.

Figure 367: Configuring a Network Interface for RIP

The screenshot shows a web interface for configuring RIP on a network interface. The breadcrumb path is "Routing Protocol > RIP > Interface". The "Action" dropdown is set to "Add". The configuration fields are as follows:

VLAN ID (1-4093)	1
Send Version	RIPv1 Compatible
Receive Version	RIPv1 and RIPv2
Authentication Type	Simple Password
Authentication Key	mighty
Instability Prevention	Poison Reverse

At the bottom right, there are "Apply" and "Revert" buttons.

To show the network interface settings configured for RIP:

1. Click Routing Protocol, RIP, Interface.
2. Select Show from the Action list.

Figure 368: Showing RIP Network Interface Settings

Routing Protocol > RIP > Interface

Action: Show

Interface Settings List Total: 2

<input type="checkbox"/>	VLAN ID	Send Version	Receive Version	Authentication Type	Authentication Key	Instability Prevention
<input type="checkbox"/>	1	RIPv1 Compatible	RIPv1 and RIPv2	Simple Password	mighty	Poison Reverse
<input type="checkbox"/>	2	RIPv1 Compatible	RIPv1 and RIPv2	No Authentication		Poison Reverse

Buttons: Delete, Revert

Displaying RIP Interface Settings Use the Routing Protocol > RIP > Statistics (Show Interface Information) page to display information about RIP interface configuration settings.

Parameters

These parameters are displayed:

- ◆ **Interface** – Source IP address of RIP router interface.
- ◆ **Auth Type** – The type of authentication used for exchanging RIPv2 protocol messages.
- ◆ **Send Version** – The RIP version to sent on this interface.
- ◆ **Receive Version** – The RIP version accepted on this interface.
- ◆ **Rcv Bad Packets** – Number of bad RIP packets received.
- ◆ **Rcv Bad Routes** – Number of bad routes received.
- ◆ **Send Updates** – Number of route changes.

Web Interface

To display RIP interface configuration settings:

1. Click Routing Protocol, RIP, Statistics.
2. Select Show Interface Information from the Action list.

Figure 369: Showing RIP Interface Settings

Routing Protocol > RIP > Statistics

Action: Show Interface Information

Interface Information Total: 3

Interface	Auth Type	Send Version	Receive Version	Rcv Bad Packets	Rcv Bad Routes	Send Updates
1.2.3.4	No Authentication	Do Not Send	RIPv1 and RIPv2	10	2	124
10.1.0.1	Simple Password	RIPv1	Do Not Receive	3	4	23
140.113.1.3	MD5	RIPv1 Compatible	RIPv2	5	5	65

Displaying Peer Router Information Use the Routing Protocol > RIP > Statistics (Show Peer Information) page to display information on neighboring RIP routers.

Parameters

These parameters are displayed:

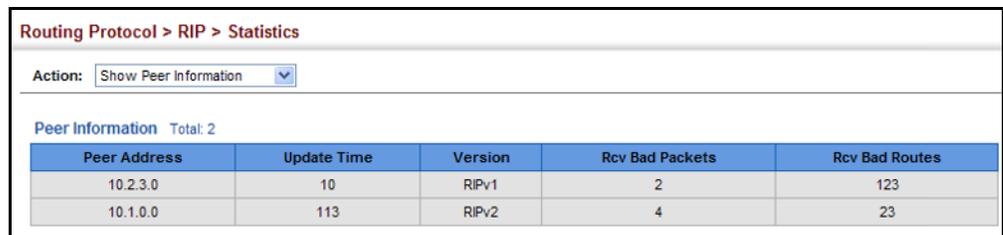
- ◆ **Peer Address** – IP address of a neighboring RIP router.
- ◆ **Update Time** – Last time a route update was received from this peer.
- ◆ **Version** – Shows whether RIPv1 or RIPv2 packets were received from this peer.
- ◆ **Rcv Bad Packets** – Number of bad RIP packets received from this peer.
- ◆ **Rcv Bad Routes** – Number of bad routes received from this peer.

Web Interface

To display information on neighboring RIP routers:

1. Click Routing Protocol, RIP, Statistics.
2. Select Show Peer Information from the Action list.

Figure 370: Showing RIP Peer Information



The screenshot shows the 'Routing Protocol > RIP > Statistics' page. At the top, there is a breadcrumb trail and an 'Action:' dropdown menu set to 'Show Peer Information'. Below this, the 'Peer Information' section is displayed with a 'Total: 2' count. A table with five columns is shown: Peer Address, Update Time, Version, Rcv Bad Packets, and Rcv Bad Routes. The table contains two rows of data.

Peer Address	Update Time	Version	Rcv Bad Packets	Rcv Bad Routes
10.2.3.0	10	RIPv1	2	123
10.1.0.0	113	RIPv2	4	23

Resetting RIP Statistics Use the Routing Protocol > RIP > Statistics (Reset Statistics) page to reset all statistics for RIP protocol messages.

Web Interface

To reset RIP statistics:

1. Click Routing Protocol, RIP, Statistics.
2. Select Reset Statistics from the Action list.
3. Click Reset.

Figure 371: Resetting RIP Statistics



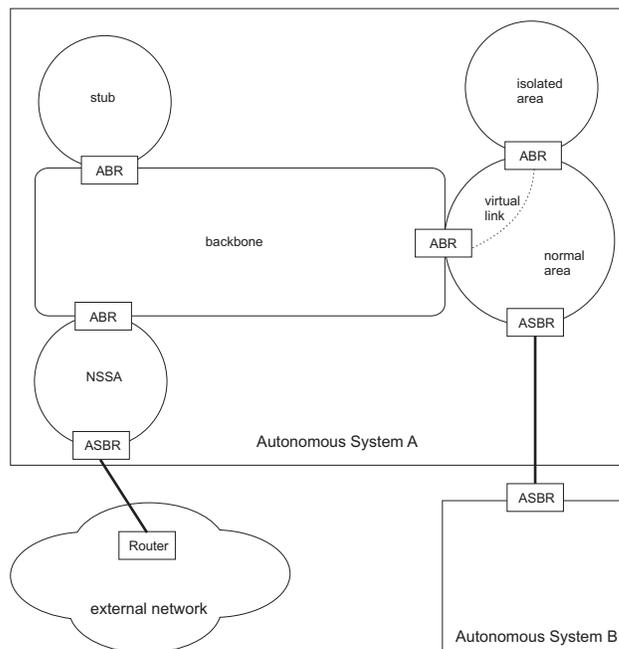
Configuring the Open Shortest Path First Protocol (Version 2)

Open Shortest Path First (OSPF) is more suited for large area networks which experience frequent changes in the links. It also handles subnets much better than RIP. OSPF protocol actively tests the status of each link to its neighbors to generate a shortest path tree, and builds a routing table based on this information. OSPF then utilizes IP multicast to propagate routing information. A separate routing area scheme is also used to further reduce the amount of routing traffic.



Note: The OSPF protocol implemented in this device is based on RFC 2328 (Version 2). It also supports RFC 1583 (early Version 2) compatibility mode to ensure that the same method is used to calculate summary route costs throughout the network when older OSPF routers exist; as well as the not-so-stubby area option (RFC 3101).

Figure 372: Configuring OSPF



Command Usage

- ◆ OSPF looks at more than just the simple hop count. When adding the shortest path to any node into the tree, the optimal path is chosen on the basis of delay,

throughput and connectivity. OSPF utilizes IP multicast to reduce the amount of routing traffic required when sending or receiving routing path updates. The separate routing area scheme used by OSPF further reduces the amount of routing traffic, and thus inherently provides another level of routing protection. In addition, all routing protocol exchanges can be authenticated. Finally, the OSPF algorithms have been tailored for efficient operation in TCP/IP Internets.

- ◆ OSPFv2 is a compatible upgrade to OSPF. It involves enhancements to protocol message authentication, and the addition of a point-to-multipoint interface which allows OSPF to run over non-broadcast networks, as well as support for overlapping area ranges.
- ◆ When using OSPF, you must organize your network (i.e., autonomous system) into normal, stub, or not-so-stubby areas; configure the ranges of subnet addresses that can be aggregated by link state advertisements; and configure virtual links for areas that do not have direct physical access to the OSPF backbone.
 - To implement OSPF for a large network, you must first organize the network into logical areas to limit the number of OSPF routers that actively exchange Link State Advertisements (LSAs). You can then define an OSPF interface by assigning an IP interface configured on this router to one of these areas. This OSPF interface will send and receive OSPF traffic to neighboring OSPF routers.
 - You can further optimize the exchange of OSPF traffic by specifying an area range that covers a large number of subnetwork addresses. This is an important technique for limiting the amount of traffic exchanged between Area Border Routers (ABRs).
 - And finally, you must specify a virtual link to any OSPF area that is not physically attached to the OSPF backbone. Virtual links can also be used to provide a redundant link between contiguous areas to prevent areas from being partitioned, or to merge backbone areas. (Note that virtual links are not supported for stubs or NSSAs.)

Defining Network Areas Based on Addresses

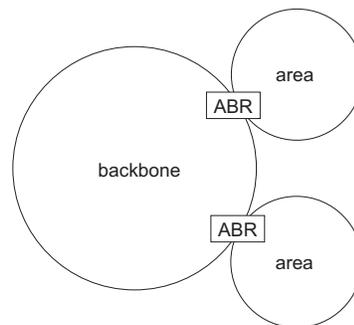
OSPF protocol broadcast messages (i.e., Link State Advertisements or LSAs) are restricted by area to limit their impact on network performance. A large network should be split up into separate OSPF areas to increase network stability, and to reduce protocol traffic by summarizing routing information into more compact messages. Each router in an area shares the same view of the network topology, including area links, route summaries for directly connected areas, and external links to other areas.

Use the Routing Protocol > OSPF > Network Area (Add) page to define an OSPF area and the interfaces that operate within this area. An autonomous system must be configured with a backbone area, designated by the area identifier 0.0.0.0. By default, all other areas are created as normal transit areas.

Routers in a normal area may import or export routing information about individual nodes. To reduce the amount of routing traffic flooded onto the network, an area can be configured to export a single summarized route that covers a broad range of network addresses within the area (page 578). To further reduce the amount of routes passed between areas, an area can be configured as a stub (page 571, page 575) or a not-so-stubby area (page 571, page 572).

Normal Area – A large OSPF domain should be broken up into several areas to increase network stability and reduce the amount of routing traffic required through the use of route summaries that aggregate a range of addresses into a single route. The backbone or any normal area can pass traffic between other areas, and are therefore known as transit areas. Each router in an area has identical routing tables. These tables may include area links, summarized links, or external links that depict the topology of the autonomous system.

Figure 373: OSPF Areas



Command Usage

- ◆ Specify an Area ID and the corresponding network address range for each OSPF broadcast area. Each area identifies a logical group of OSPF routers that actively exchange Link State Advertisements (LSAs) to ensure that they share an identical view of the network topology.
- ◆ Each area must be connected to a backbone area. This area passes routing information between other areas in the autonomous system. All routers must be connected to the backbone, either directly, or through a virtual link if a direct physical connection is not possible.
- ◆ All areas are created as normal transit areas using the Network Area (Add) page. A normal area (or transit area) can send and receive external LSAs. If necessary, an area can be configured as a not-so-stubby area (NSSA) that can import external route information into its area, or as a stubby area that cannot send or receive external LSAs.
- ◆ An area must be assigned a range of subnetwork addresses. This area and the corresponding address range forms a routing interface, and can be configured to aggregate LSAs from all of its subnetwork addresses and exchange this information with other routers in the network as described under “[Configuring Area Ranges \(Route Summarization for ABRs\)](#)” on page 578.

- ◆ If an address range overlaps other network areas, the router will use the network area with the address range that most closely matches the interface address. Also, note that if a more specific address range is removed from an area, the interface belonging to that range may still remain active if a less specific address range covering that area has been specified.

Parameters

These parameters are displayed:

- ◆ **Process ID** – Protocol identifier used to distinguish between multiple routing instances. (Range: 1-65535)
- ◆ **IP Address** – Address of the interfaces to add to the area.
- ◆ **Netmask** – Network mask of the address range to add to the area.
- ◆ **Area ID** – Area to which the specified address or range is assigned. An OSPF area identifies a group of routers that share common routing information. The area ID can be in the form of an IPv4 address, or as a four octet unsigned integer ranging from 0-4294967295.

Set the area ID to the same value for all routers on a network segment using the network mask to add one or more interfaces to an area.

Web Interface

To define an OSPF area and the interfaces that operate within this area:

1. Click Routing Protocol, OSPF, Network Area.
2. Select Add from the Action list.
3. Configure a backbone area that is contiguous with all the other areas in the network, and configure an area for all of the other OSPF interfaces.
4. Click Apply

Figure 374: Defining OSPF Network Areas Based on Addresses

The screenshot shows a web interface for configuring an OSPF network area. The breadcrumb navigation is "Routing Protocol > OSPF > Network Area". At the top, there is an "Action:" dropdown menu currently set to "Add". Below this, there are four input fields for configuration: "Process ID (1-65535)" with the value "1", "IP Address" with "192.168.0.0", "Netmask" with "255.255.255.0", and "Area ID" with "0.0.0.0". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

To to show the OSPF areas and the assigned interfaces:

1. Click Routing Protocol, OSPF, Network Area.
2. Select Show from the Action list.

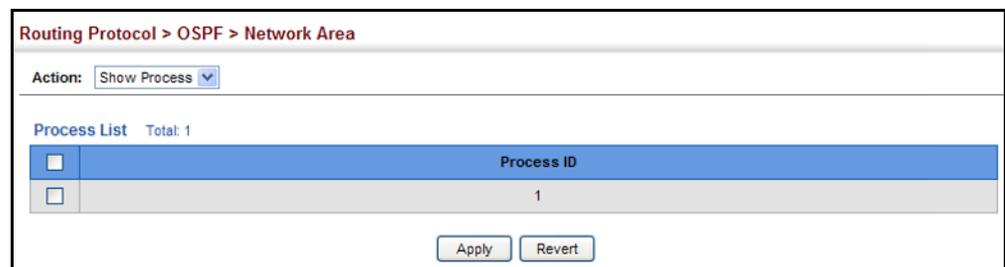
Figure 375: Showing OSPF Network Areas



To to show the OSPF process identifiers:

1. Click Routing Protocol, OSPF, Network Area.
2. Select Show Process from the Action list.

Figure 376: Showing OSPF Process Identifiers



Configuring General Protocol Settings

To implement dynamic OSPF routing, first assign VLAN groups to each IP subnet to which this router will be attached (as described in the preceding section), then use the Routing Protocol > OSPF > System (Configure) page to assign an Router ID to this device, and set the other basic protocol parameters.

Parameters

These parameters are displayed:

- ◆ **Process ID** – Protocol identifier as configured on the Routing Protocol > OSPF > Network Area (Add) page. (Range: 1-65535)

General Information

- ◆ **RFC1583 Compatible** – If one or more routers in a routing domain are using early Version 2 of OSPF, this router should use RFC 1583 (early OSPFv2)

compatibility mode to ensure that all routers are using the same RFC for calculating summary route costs. Enable this field to force the router to calculate summary route costs using RFC 1583. (Default: Disabled)

When RFC 1583 compatibility is enabled, only cost is used when choosing among multiple AS-external LSAs advertising the same destination. When disabled, preference is based on type of path, using cost only to break ties (see RFC 2328).

If there are any OSPF routers in an area exchanging summary information (specifically, ABRs) which have not been upgraded to OSPFv2 (RFC 2328), RFC 1583 should be used on the newly upgraded OSPFv2 routers to ensure compatibility with routers still running older OSPFv2 code.

- ◆ **OSPF Router ID** – Assigns a unique router ID for this device within the autonomous system for the current OSPF process.

The router ID must be unique for every router in the autonomous system. Note that the router ID cannot be set to 0.0.0.0.

If this router already has registered neighbors, the new router ID will be used when the router is rebooted, or manually restarted using the “no router ospf” command followed by the “router ospf” command.

- ◆ **Auto Cost** – Calculates the cost for an interface by dividing the reference bandwidth by the interface bandwidth. The reference bandwidth is defined in Mbits per second. (Range: 1-4294967)

By default, the cost is 0.1 for Gigabit ports, and 0.01 for 10 Gigabit ports. A higher reference bandwidth can be used for aggregate links to indicate preferred use as a lower cost interface.

- ◆ **SPF Hold Time** – The hold time between making two consecutive shortest path first (SPF) calculations. (Range: 0-65535 seconds; Default: 10 seconds)

Setting the SPF holdtime to 0 means that there is no delay between consecutive calculations.

- ◆ **SPF Delay Time** – The delay after receiving a topology change notification and starting the SPF calculation. (Range: 0-65535 seconds; Default: 5 seconds)

Using a low value for the delay and hold time allows the router to switch to a new path faster, but uses more CPU processing time.

- ◆ **Default Metric** – The default metric for external routes imported from other protocols. (Range: 0-16777214; Default: 20)

A default metric must be used to resolve the problem of redistributing external routes from other protocols that use incompatible metrics.

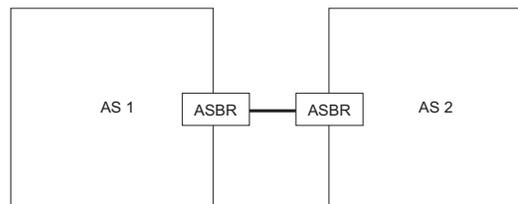
This default metric does not override the metric value set on the Redistribute configuration screen (see [page 580](#)). When a metric value has not been configured on the Redistribute page, the default metric configured on the System configuration page sets the metric value to be used for all imported external routes.

Default Information

- ◆ **Originate Default Route**¹⁷ – Generates a default external route into an autonomous system. Note that the **Advertise Default Route** field must also be properly configured. (Default: Disabled)

When this feature is used to redistribute routes into a routing domain (that is, an Autonomous System), this router automatically becomes an Autonomous System Boundary Router (ASBR). This allows the router to exchange routing information with boundary routers in other autonomous systems to which it may be attached. If a router is functioning as an ASBR, then every other router in the autonomous system can learn about external routes from this device.

Figure 377: AS Boundary Router



- ◆ **Advertise Default Route**¹⁷ – The router can advertise a default external route into the autonomous system (AS). (Options: Not Always, Always; Default: Not Always)
 - **Always** – The router will advertise itself as a default external route for the local AS, even if a default external route does not actually exist. (To define a default route, see [“Configuring Static Routes” on page 526.](#))
 - **NotAlways** – It can only advertise a default external route into the AS if it has been configured to import external routes through RIP or static routes, and such a route is known. (See [“Redistributing External Routes” on page 580.](#))
- ◆ **External Metric Type**¹⁷ – The external link type used to advertise the default route. Type 1 route advertisements add the internal cost to the external route metric. Type 2 routes do not add the internal cost metric. When comparing Type 2 routes, the internal cost is only used as a tie-breaker if several Type 2 routes have the same cost. (Default: Type 2)
- ◆ **Default External Metric**¹⁷ – Metric assigned to the default route. (Range: 0-16777215; Default: 20)

The metric for the default external route is used to calculate the path cost for traffic passed from other routers within the AS out through the ASBR.

Redistribution of routing information from other protocols is controlled by the Redistribute function (see [page 580](#)).

¹⁷. These are configured with the “default-information originate” command.

Web Interface

To configure general settings for OSPF:

1. Click Routing Protocol, OSPF, System.
2. Select Configure from the Action list.
3. Select a Process ID, and then specify the Router ID and other global attributes as required. For example, by setting the Auto Cost to 10000, the cost of using an interface is set to 10 for Gigabit ports, and 1 for 10 Gigabit ports.
4. Click Apply

Figure 378: Configure General Settings for OSPF

Displaying Administrative Settings and Statistics

Use the Routing Protocol > OSPF > System (Show) page to display general administrative settings and statistics for OSPF.

Parameters

These parameters are displayed:

Table 44: OSPF System Information

Parameter	Description
Router ID Type	Indicates if the router ID was manually configured or automatically generated by the system.
Rx LSAs	The number of link-state advertisements that have been received.

Table 44: OSPF System Information (Continued)

Parameter	Description
Originate LSAs	The number of new link-state advertisements that have been originated.
AS LSA Count	The number of autonomous system LSAs in the link-state database.
External LSA Count	The number of external link-state advertisements in the link-state database.
External LSA Checksum	Checksum of the external link-state advertisement database.
Admin Status	Indicates if there are one or more configured OSPF areas with an active interface (that is, a Layer 3 interface that is enabled and up).
ABR Status (Area Border Router)	Indicates if this router connects directly to networks in two or more areas. An area border router runs a separate copy of the Shortest Path First algorithm, maintaining a separate routing database for each area.
ASBR Status (Autonomous System Boundary Router)	Indicates if this router exchanges routing information with boundary routers in other autonomous systems to which it may be attached. If a router is enabled as an ASBR, then every other router in the autonomous system can learn about external routes from this device.
Restart Status	Indicates if the OSPF process is in graceful-restart state.
Area Number	The number of configured areas attached to this router.
Version Number	The OSPF version number. The OSPF protocol implemented in this device is based on RFC 2328 (Version 2). It also supports RFC 1583 (early Version 2) compatibility mode.

Web Interface

To show administrative settings and statistics for OSPF:

To display general settings for OSPF:

1. Click Routing Protocol, OSPF, System.
2. Select Show from the Action list.
3. Select a Process ID.

Figure 379: Showing General Settings for OSPF

Routing Protocol > OSPF > System			
Action: <input type="button" value="Show"/>			
Process ID: <input type="button" value="1"/>			
OSPF System Information			
Router ID Type	Auto	Admin Status	Enabled
Rx LSAs	5	ABR Status	Disabled
Originate LSAs	5	ASBR Status	Disabled
AS LSA Count	0	Restart Status	Disabled
External LSA Count	1	Area Number	2
External LSA Checksum	CD97	Version Number	2

Adding an NSSA or Stub Use the Routing Protocol > OSPF > Area (Configure Area – Add Area) page to add a not-so-stubby area (NSSA) or a stubby area (Stub).

Command Usage

- ◆ This router supports up to 5 stubs or NSSAs.

Parameters

These parameters are displayed:

- ◆ **Process ID** – Protocol identifier as configured on the Routing Protocol > OSPF > Network Area (Add) page. (Range: 1-65535)
- ◆ **Area ID** – Identifier for a not-so-stubby area (NSSA) or stub. The area ID can be in the form of an IPv4 address, or as a four octet unsigned integer ranging from 0-4294967295.

Set the area ID to the same value for all routers on a network segment using the network mask to add one or more interfaces to an area.
- ◆ **Area Type** – Specifies an NSSA or stub.

Web Interface

To add an NSSA or stub to the OSPF administrative domain:

1. Click Routing Protocol, OSPF, Area.
2. Select Configure Area from the Step list.
3. Select Add Area from the Action list.
4. Select a Process ID, enter the area identifier, and set the area type to NSSA or Stub.
5. Click Apply

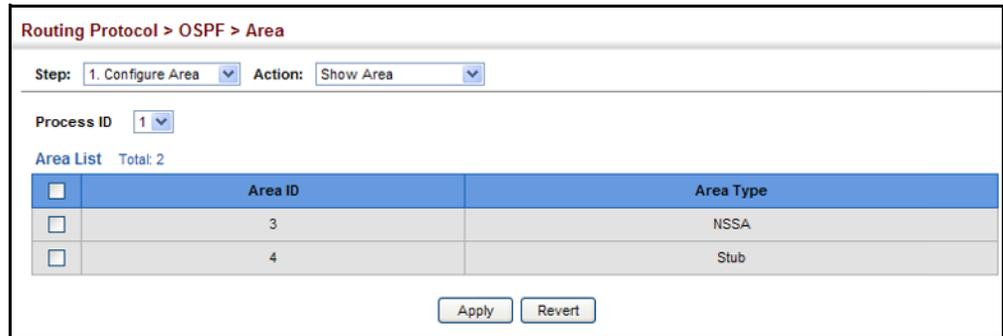
Figure 380: Adding an NSSA or Stub

The screenshot shows a web interface for configuring OSPF areas. At the top, the breadcrumb is "Routing Protocol > OSPF > Area". Below this, there are two dropdown menus: "Step:" with "1. Configure Area" selected, and "Action:" with "Add Area" selected. The main configuration area contains three fields: "Process ID" with a dropdown menu showing "1", "Area ID" with a text input field containing "3", and "Area Type" with a dropdown menu showing "NSSA". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

To show the NSSA or stubs added to the specified OSPF domain:

1. Click Routing Protocol, OSPF, Area.
2. Select Configure Area from the Step list.
3. Select Show Area from the Action list.
4. Select a Process ID.

Figure 381: Showing NSSAs or Stubs



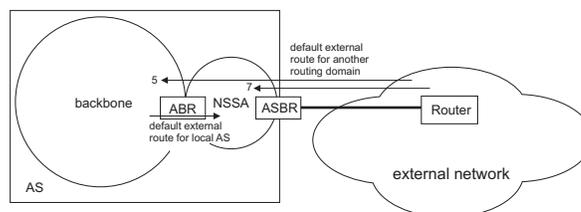
Configuring NSSA Settings Use the Routing Protocol > OSPF > Area (Configure Area – Configure NSSA Area) page to configure protocol settings for a not-so-stubby area (NSSA).

An NSSA can be configured to control the use of default routes for Area Border Routers (ABRs) and Autonomous System Boundary Routers (ASBRs), or external routes learned from other routing domains and imported through an ABR.

An NSSA is similar to a stub. It blocks most external routing information, and can be configured to advertise a single default route for traffic passing between the NSSA and other areas within the autonomous system (AS) when the router is an ABR.

An NSSA can also import external routes from one or more small routing domains that are not part of the AS, such as a RIP domain or locally configured static routes. This external AS routing information is generated by the NSSA's ASBR and advertised only within the NSSA. By default, these routes are not flooded onto the backbone or into any other area by ABRs. However, the NSSA's ABRs will convert NSSA external LSAs (Type 7) into external LSAs (Type-5) which are propagated into other areas within the AS.

Figure 382: OSPF NSSA



Command Usage

- ◆ Before creating an NSSA, first specify the address range for the area (see [“Defining Network Areas Based on Addresses” on page 563](#)). Then create an NSSA as described under [“Adding an NSSA or Stub” on page 571](#).
- ◆ NSSAs cannot be used as a transit area, and should therefore be placed at the edge of the routing domain.
- ◆ An NSSA can have multiple ABRs or exit points. However, all of the exit points and local routers must contain the same external routing data so that the exit point does not need to be determined for each external destination.
- ◆ There are no external routes in an OSPF stub area, so routes cannot be redistributed from another protocol into a stub area. On the other hand, an NSSA allows external routes from another protocol to be redistributed into its own area, and then leaked to adjacent areas.
- ◆ Routes that can be advertised with NSSA external LSAs include network destinations outside the AS learned through OSPF, the default route, static routes, routes derived from other routing protocols such as RIP, or directly connected networks that are not running OSPF.
- ◆ An NSSA can be used to simplify administration when connecting a central site using OSPF to a remote site that is using a different routing protocol. OSPF can be easily extended to cover the remote connection by defining the area between the central router and the remote router as an NSSA.

Parameters

These parameters are displayed:

- ◆ **Process ID** – Process ID as configured in the Network Area configuration screen (see [page 563](#)).
- ◆ **Area ID** – Identifier for a not-so-stubby area (NSSA).
- ◆ **Translator Role** – Indicates NSSA-ABR translator role for converting Type 7 external LSAs into Type 5 external LSAs. These roles include:
 - **Never** – A router that never translates NSSA LSAs to Type-5 external LSAs.
 - **Always** – A router that always translates NSSA LSA to Type-5 external LSA.
 - **Candidate** – A router translates NSSA LSAs to Type-5 external LSAs if elected.
- ◆ **Redistribute** – Disable this option when the router is an NSSA Area Border Router (ABR) and routes only need to be imported into normal areas (see [“Redistributing External Routes” on page 580](#)), but not into the NSSA. In other words, redistribution should be disabled to prevent the NSSA ABR from advertising external routing information (learned through routers in other areas) into the NSSA. (Default: Enabled)

- ◆ **Originate Default Information** – When the router is an NSSA Area Border Router (ABR) or an NSSA Autonomous System Boundary Router (ASBR), this option causes it to generate a Type-7 default LSA into the NSSA. This default provides a route to other areas within the AS for an NSSA ABR, or to areas outside the AS for an NSSA ASBR. (Default: Disabled)

An NSSA is similar to a stub, because when the router is an ABR, it can send a default route for other areas in the AS into the NSSA using the Originate Default Information option. However, an NSSA is different from a stub, because when the router is an ASBR, it can import a default external AS route (for routing protocol domains adjacent to the NSSA but not within the OSPF AS) into the NSSA using this option.

- ◆ **Metric Type** – Type 1 or Type 2 external routes. When using Type 2, routers do not add internal cost to the external route metric. (Default: Type 2)
- ◆ **Metric** – Metric assigned to Type-7 default LSAs. (Range: 0-16777214; Default: 1)
- ◆ **Default Cost** – Cost for the default summary route sent into an NSSA from an area border router (ABR). (Range: 0-16777215; Default: 0)

Note that when the default cost is set to “0,” the router will not advertise a default route into the attached NSSA.

- ◆ **Summary** – Controls the use of summary routes. (Default: Summary)
 - **Summary** – Unlike stub areas, all Type-3 summary LSAs will be imported into NSSAs to ensure that internal routes are always chosen over Type-7 NSSA external routes.
 - **No Summary** – Allows an area to retain standard NSSA features, but does not inject inter-area routes (Type-3 and Type-4 summary routes) into this area. Instead, it advertises a default route as a Type-3 LSA.

Web Interface

To configure protocol settings for an NSSA:

1. Click Routing Protocol, OSPF, Area.
2. Select Configure Area from the Step list.
3. Select Configure NSSA Area from the Action list.
4. Select a Process ID, and modify the routing behavior for an NSSA.
5. Click Apply

Figure 383: Configuring Protocol Settings for an NSSA

Routing Protocol > OSPF > Area

Step: 1. Configure Area Action: Configure NSSA Area

Process ID: 1

NSSA Area List Total: 1

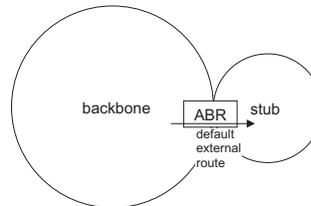
Area ID	Translator Role	Redistribute	Originate Default Information	Metric Type	Metric (0-16777215)	Default Cost (0-16777215)	Summary
192.168.2.0	Candidate	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	2	34	0	Summary

Apply Revert

Configuring Stub Settings Use the Routing Protocol > OSPF > Area (Configure Area – Configure Stub Area) page to configure protocol settings for a stub.

A stub does not accept external routing information. Instead, an area border router adjacent to a stub can be configured to send a default external route into the stub for all destinations outside the local area or the autonomous system. This route will also be advertised as a single entry point for traffic entering the stub. Using a stub can significantly reduce the amount of topology data that has to be exchanged over the network.

Figure 384: OSPF Stub Area



By default, a stub can only pass traffic to other areas in the autonomous system through the default external route. However, an area border router can also be configured to send Type 3 summary link advertisements into the stub about subnetworks located elsewhere in the autonomous system.

Command Usage

- ◆ Before creating a stub, first specify the address range for the area (see [“Defining Network Areas Based on Addresses”](#) on page 563). Then create a stub as described under [“Adding an NSSA or Stub”](#) on page 571.
- ◆ Stubs cannot be used as a transit area, and should therefore be placed at the edge of the routing domain.
- ◆ A stub can have multiple ABRs or exit points. However, all of the exit points and local routers must contain the same external routing data so that the exit point does not need to be determined for each external destination.

Parameters

These parameters are displayed:

- ◆ **Process ID** – Process ID as configured in the Network Area configuration screen (see [page 563](#)).
- ◆ **Area ID** – Identifier for a stub.
- ◆ **Default Cost** – Cost for the default summary route sent into a stub from an area border router (ABR). (Range: 0-16777215; Default: 0)

Note that if the default cost is set to “0,” the router will not advertise a default route into the attached stub.

- ◆ **Summary** – Controls the use of summary routes.
 - **Summary** – Allows an Area Border Router (ABR) to send a summary link advertisement into the stub area.
 - **No Summary** – Stops an ABR from sending a summary link advertisement into a stub area.

Routing table space is saved in a stub by blocking Type-4 AS summary LSAs and Type 5 external LSAs. This option can be used to completely isolate the stub by also stopping an ABR from sending Type-3 summary LSAs that advertise the default route for destinations external to the local area or the autonomous system.

Define an area as a totally stubby area only if routers in the area do not require summary LSAs from other areas.

Web Interface

To configure protocol settings for a stub:

1. Click Routing Protocol, OSPF, Area.
2. Select Configure Area from the Step list.
3. Select Configure Stub Area from the Action list.
4. Select a Process ID, and modify the routing behavior for a stub.
5. Click Apply

Figure 385: Configuring Protocol Settings for a Stub

The screenshot shows the 'Routing Protocol > OSPF > Area' configuration page. At the top, there are two dropdown menus: 'Step: 1. Configure Area' and 'Action: Configure Stub Area'. Below this, there is a 'Process ID' dropdown menu set to '1'. The main section is titled 'Stub Area List' with a 'Total: 1' indicator. It contains a table with three columns: 'Area ID', 'Default Cost (0-16777215)', and 'Summary'. The table has one row with the following data: Area ID '192.168.3.0', Default Cost '1', and Summary 'Summary'. Below the table are 'Apply' and 'Revert' buttons.

Area ID	Default Cost (0-16777215)	Summary
192.168.3.0	1	Summary

Displaying Information on NSSA and Stub Areas

Use the Routing Protocol > OSPF > Area (Show Information) page to protocol information on NSSA and Stub areas.

Parameters

These parameters are displayed:

- ◆ **Process ID** – Process ID as configured in the Network Area configuration screen (see [page 563](#)).
- ◆ **Area ID** – Identifier for a not-so-stubby area (NSSA) or stub.
- ◆ **SPF Runs** – The number of times the Shortest Path First algorithm has been run for this area.
- ◆ **ABR Count** – The number of Area Border Routers attached to this area.
- ◆ **ASBR Count** – The number of Autonomous System Boundary Routers attached to this area.
- ◆ **LSA Count** – The number of new link-state advertisements that have been originated.
- ◆ **LSA Checksum Sum** – The sum of the link-state advertisements' LS checksums contained in this area's link-state database.

Web Interface

To display information on NSSA and stub areas:

1. Click Routing Protocol, OSPF, Area.
2. Select Show Information from the Action list.
3. Select a Process ID.

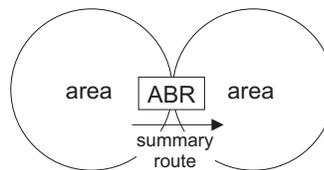
Figure 386: Displaying Information on NSSA and Stub Areas

Routing Protocol > OSPF > Area					
Step: 2. Show Information					
Process ID: 1					
Area Information List Total: 4					
Area ID	SPF Runs	ABR Count	ASBR Count	LSA Count	LSA Checksum Sum
0.0.0.1	0	0	0	0	0
0.0.0.2	0	0	0	0	0
0.0.0.3	10	10	10	10	10
0.0.0.4	0	0	0	0	0

Configuring Area Ranges (Route Summarization for ABRs)

An OSPF area can include a large number of nodes. If the Area Border Router (ABR) has to advertise route information for each of these nodes, this wastes a lot of bandwidth and processor time. Instead, you can use the Routing Protocol > OSPF > Area Range (Add) page to configure an ABR to advertise a single summary route that covers all the individual networks within its area. When using route summaries, local changes do not have to be propagated to other area routers. This allows OSPF to be easily scaled for larger networks, and provides a more stable network topology.

Figure 387: Route Summarization for ABRs



Command Usage

- ◆ Use the Area Range configuration page to summarize intra-area routes, and advertise this information to other areas through Area Border Routers (ABRs). The summary route for an area is defined by an IP address and network mask. You therefore need to structure each area with a contiguous set of addresses so that all routes in the area fall within an easily specified range. If it is not possible to use one contiguous set of addresses, then the routes can be summarized for several area ranges. This router also supports Variable Length Subnet Masks (VLSMs), so you can summarize an address range on any bit boundary in a network address.
- ◆ To summarize the external LSAs imported into your autonomous system (i.e., local routing domain), use the Summary Address configuration screen ([page 578](#)).
- ◆ This router supports up five summary routes for area ranges.

Parameters

These parameters are displayed:

- ◆ **Process ID** – Process ID as configured in the Network Area configuration screen (see [page 563](#)).
- ◆ **Area ID** – Identifies an area for which the routes are summarized. The area ID can be in the form of an IPv4 address, or also as a four octet unsigned integer ranging from 0-4294967295.
- ◆ **Range Network** – Base address for the routes to summarize.
- ◆ **Range Netmask** – Network mask for the summary route.
- ◆ **Advertising** – Indicates whether or not to advertise the summary route. If the routes are set to be advertised, the router will issue a Type 3 summary LSA for each specified address range. If the summary is not advertised, the specified routes remain hidden from the rest of the network. (Default: Advertise)

Web Interface

To configure a route summary for an area range:

1. Click Routing Protocol, OSPF, Area Range.
2. Select Add from the Action list.
3. Specify the process ID, area identifier, the base address and network mask, and select whether or not to advertise the summary route to other areas.
4. Click Apply

Figure 388: Configuring Route Summaries for an Area Range

Routing Protocol > OSPF > Area Range

Action: Add

Process ID: 1

Area ID: 192.168.0.0

Range Network: 192.168.0.0

Range Netmask: 255.255.0.0

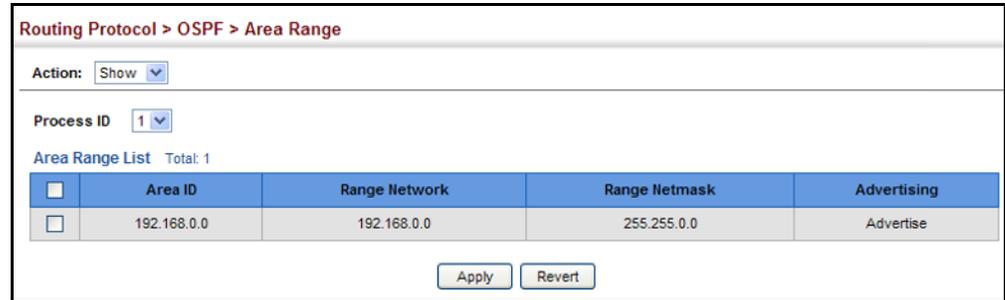
Advertising: Advertise

Apply Revert

To show the configured route summaries:

1. Click Routing Protocol, OSPF, Area Range.
2. Select Show from the Action list.
3. Select the process ID.

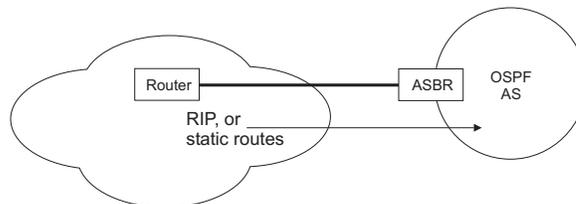
Figure 389: Showing Configured Route Summaries



Redistributing External Routes

Use the Routing Protocol > OSPF > Redistribute (Add) page to import external routing information from other routing protocols, static routes, or directly connected routes into the autonomous system, and to generate AS-external-LSAs.

Figure 390: Redistributing External Routes



Command Usage

- ◆ This router supports redistribution for all currently connected routes, entries learned through RIP, and static routes.
- ◆ When you redistribute external routes into an OSPF autonomous system (AS), the router automatically becomes an autonomous system boundary router (ASBR).
- ◆ However, if the router has been configured as an ASBR via the General Configuration screen, but redistribution is not enabled, the router will only generate a “default” external route into the AS if it has been configured to “always” advertise a default route even if an external route does not actually exist (page 566).

Parameters

These parameters are displayed:

- ◆ **Process ID** – Process ID as configured in the Network Area configuration screen (see [page 563](#)).
- ◆ **Protocol Type** – Specifies the external routing protocol type for which routing information is to be redistributed into the local routing domain. (Options: RIP, Static, Connected, BGP; Default: RIP)
- ◆ **Metric Type** – Indicates the method used to calculate external route costs. (Options: Type 1, Type 2; Default: Type 1)

Metric type specifies the way to advertise routes to destinations outside the autonomous system (AS) through External LSAs. Specify Type 1 to add the internal cost metric to the external route metric. In other words, the cost of the route from any router within the AS is equal to the cost associated with reaching the advertising ASBR, plus the cost of the external route. Specify Type 2 to only advertise the external route metric.

- ◆ **Metric** – Metric assigned to all external routes for the specified protocol. (Range: 1-65535; Default: 10)
The metric value specified for redistributed routes supersedes the Default External Metric specified in the Routing Protocol > OSPF > System screen ([page 566](#)).
- ◆ **Tag** – A tag placed in the AS-external LSA to identify a specific external routing domain, or to pass additional information between routers. (Range: 0-4294967295)
A tag can be used to distinguish between routes learned from different external autonomous systems (other routing protocols). For example, if there are two ASBRs in a routing domain: A and B. ASBR A can be configured to redistribute routes learned from RIP domain 1 (identified by tag 1) and ASBR B can redistribute routes learned from RIP domain 2 (identified by tag 2).

Web Interface

To configure the router to import external routing information:

1. Click Routing Protocol, OSPF, Redistribute.
2. Select Add from the Action list.
3. Specify the process ID, the protocol type to import, the metric type, path cost, and optional tag.
4. Click Apply.

Figure 391: Importing External Routes

To show the imported external route types:

1. Click Routing Protocol, OSPF, Redistribute.
2. Select Show from the Action list.
3. Select the process ID.

Figure 392: Showing Imported External Route Types

	Protocol Type	Metric Type	Metric	Tag
<input type="checkbox"/>	Static	1	2	3

Configuring Summary Addresses (for External AS Routes)

Redistributing routes from other protocols into OSPF normally requires the router to advertise each route individually in an external LSA as described in the preceding section. To reduce the number of protocol messages required to redistribute these external routes, an Autonomous System Boundary Router (ASBR) can instead be configured to redistribute routes learned from other protocols into all attached autonomous systems.

To reduce the amount of external LSAs sent to other autonomous systems, you can use the Routing Protocol > OSPF > Summary Address (Add) page to configure the router to advertise an aggregate route that consolidates a broad range of external addresses. This helps both to decrease the number of external LSAs advertised and the size of the OSPF link state database.

Command Usage

- ◆ If you are not sure what address ranges to consolidate, first enable external route redistribution via the Redistribute configuration screen, view the routes

imported into the routing table, and then configure one or more summary addresses to reduce the size of the routing table and consolidate these external routes for advertising into the local domain.

- ◆ To summarize routes sent between OSPF areas, use the Area Range Configuration screen ([page 578](#)).
- ◆ This router supports up to 20 Type-5 summary routes.

Parameters

These parameters are displayed:

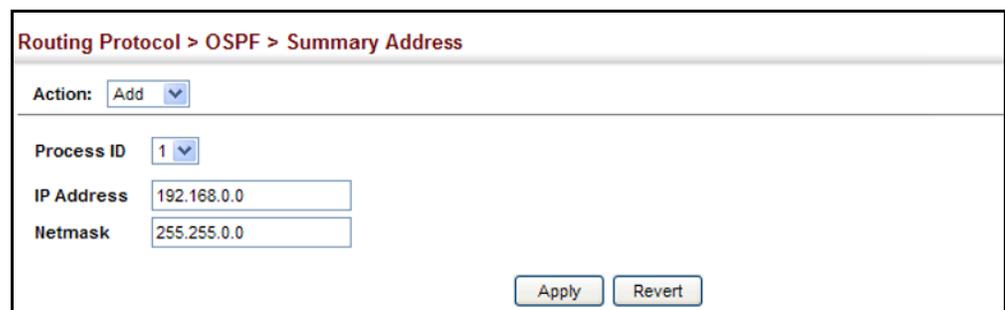
- ◆ **Process ID** – Process ID as configured in the Network Area configuration screen (see [page 563](#)).
- ◆ **IP Address** – Summary address covering a range of addresses.
- ◆ **Netmask** – Network mask for the summary route.

Web Interface

To configure the router to summarize external routing information:

1. Click Routing Protocol, OSPF, Summary Address.
2. Select Add from the Action list.
3. Specify the process ID, the base address and network mask.
4. Click Apply.

Figure 393: Summarizing External Routes



Routing Protocol > OSPF > Summary Address

Action: Add ▾

Process ID 1 ▾

IP Address 192.168.0.0

Netmask 255.255.0.0

Apply Revert

To show the summary addresses for external routes:

1. Click Routing Protocol, OSPF, Summary Address.
2. Select Show from the Action list.
3. Select the process ID.

Figure 394: Showing Summary Addresses for External Routes

Routing Protocol > OSPF > Summary Address		
Action:	Show	
Process ID	1	
Summary Address List Total: 1		
<input type="checkbox"/>	IP Address	Netmask
<input type="checkbox"/>	192.168.0.0	255.255.0.0
<input type="button" value="Apply"/> <input type="button" value="Revert"/>		

Configuring OSPF Interfaces

You should specify a routing interface for any local subnet that needs to communicate with other network segments located on this router or elsewhere in the network. First configure a VLAN for each subnet that will be directly connected to this router, assign IP interfaces to each VLAN (i.e., one primary interface and one or more secondary interfaces), and then use the Network Area configuration page to assign an interface address range to an OSPF area.

After assigning a routing interface to an OSPF area, use the Routing Protocol > OSPF > Interface (Configure by VLAN) or (Configure by Address) page to configure the interface-specific parameters used by OSPF to set the cost used to select preferred paths, select the designated router, control the timing of link state advertisements, and specify the method used to authenticate routing messages.

Command Usage

- ◆ The Configure by VLAN page is used to set the OSPF interface settings for the all areas assigned to a VLAN on the Network Area (Add) page (see [page 563](#)).
- ◆ The Configure by Address page is used to set the OSPF interface settings for a specific area assigned to a VLAN on the Network Area (Add) page (see [page 563](#)).

Parameters

These parameters are displayed:

- ◆ **VLAN ID** – A VLAN to which an IP interface has been assigned.
- ◆ **IP Address** – Address of the interfaces assigned to a VLAN on the Network Area (Add) page.
This parameter only applies to the Configure by Address page.
- ◆ **Cost** – Sets the cost of sending a protocol packet on an interface, where higher values indicate slower ports. (Range: 1-65535; Default: 1)

The interface cost indicates the overhead required to send packets across a certain interface. This is advertised as the link cost in router link state advertisements.

Routes are assigned a metric equal to the sum of all metrics for each interface link in the route.

This router uses a default cost of 1 for all ports. Therefore, if you install a 10 Gigabit module, you need to reset the cost for all of the 1 Gbps ports to a value greater than 1 to reflect the actual interface bandwidth.

- ◆ **Router Priority** – Sets the interface priority for this router. (Range: 0-255; Default: 1)

This priority determines the designated router (DR) and backup designated router (BDR) for each OSPF area. The DR forms an active adjacency to all other routers in the area to exchange routing topology information. If for any reason the DR fails, the BDR takes over this role.

Set the priority to zero to prevent a router from being elected as a DR or BDR. If set to any value other than zero, the router with the highest priority becomes the DR and the router with the next highest priority becomes the BDR. If two or more routers are set to the same highest priority, the router with the higher ID will be elected.

If a DR already exists for an area when this interface comes up, the new router will accept the current DR regardless of its own priority. The DR will not change until the next time the election process is initiated.

Configure router priority for multi-access networks only and not for point-to-point networks.

- ◆ **Hello Interval** – Sets the interval between sending hello packets on an interface. This interval must be set to the same value for all routers on the network. (Range: 1-65535 seconds; Default: 10)

Hello packets are used to inform other routers that the sending router is still active. Setting the hello interval to a smaller value can reduce the delay in detecting topological changes, but will increase routing traffic.

- ◆ **Dead Interval** – Sets the interval at which hello packets are not seen before neighbors declare the router down. This interval must be set to the same value for all routers on the network. (Range: 1-65535 seconds; Default: 40, or 4 times the Hello Interval)

The dead-interval is advertised in the router's hello packets. It must be a multiple of hello-interval and be the same for all routers on a specific network.

- ◆ **Transmit Delay** – Sets the estimated time to send a link-state update packet over an interface. (Range: 1-65535 seconds; Default: 1 second)

LSAs have their age incremented by this delay before transmission. You should consider both the transmission and propagation delays for an interface when estimating this delay. Set the transmit delay according to link speed, using larger values for lower-speed links.

If this delay is not added, the time required to transmit an LSA over the link is not taken into consideration by the routing process. On slow links, the router may send packets more quickly than devices can receive them. To avoid this

problem, you can use the transmit delay to force the router to wait a specified interval between transmissions.

- ◆ **Retransmit Interval** – Sets the time between re-sending link-state advertisements. (Range: 1-65535 seconds; Default: 5 seconds)

A router will resend an LSA to a neighbor if it receives no acknowledgment after the specified retransmit interval. The retransmit interval should be set to a conservative value that provides an adequate flow of routing information, but does not produce unnecessary protocol traffic. Note that this value should be larger for virtual links.

Set this interval to a value that is greater than the round-trip delay between any two routers on the attached network to avoid unnecessary retransmissions.

- ◆ **Authentication Type** – Specifies the authentication type used for an interface. (Options: None, Simple, MD5; Default: None)

Use authentication to prevent routers from inadvertently joining an unauthorized area. Configure routers in the same area with the same password (or key). All neighboring routers on the same network with the same password will exchange routing data.

When using simple password authentication, a password is included in the packet. If it does not match the password configured on the receiving router, the packet is discarded. This method provides very little security as it is possible to learn the authentication key by snooping on routing protocol packets.

When using Message-Digest 5 (MD5) authentication, the router uses the MD5 algorithm to verify data integrity by creating a 128-bit message digest from the authentication key. Without the proper key and key-id, it is nearly impossible to produce any message that matches the prespecified target message digest.

The Message Digest Key ID and Authentication Key and must be used consistently throughout the autonomous system.

- ◆ **Authentication Key** – Assign a plain-text password used by neighboring routers to verify the authenticity of routing protocol messages. (Range: 1-8 characters for simple password or 1-16 characters for MD5 authentication; Default: no key)

When plain-text or Message-Digest 5 (MD5) authentication is enabled as described in the preceding item, this password (key) is inserted into the OSPF header when routing protocol packets are originated by this device.

A different password can be assigned to each network interface, but the password must be used consistently on all neighboring routers throughout a network (that is, autonomous system). All neighboring routers in the same network with the same password will exchange routing data.

- ◆ **Message Digest Key ID** – Assigns a key identifier used in conjunction with the authentication key to verify the authenticity of routing protocol messages sent to neighboring routers. (Range: 1-255; Default: none)

Normally, only one key is used per interface to generate authentication information for outbound packets and to authenticate incoming packets. Neighbor routers must use the same key identifier and key value.

When changing to a new key, the router will send multiple copies of all protocol messages, one with the old key and another with the new key. Once all the neighboring routers start sending protocol messages back to this router with the new key, the router will stop using the old key. This rollover process gives the network administrator time to update all of the routers on the network without affecting the network connectivity. Once all the network routers have been updated with the new key, the old key should be removed for security reasons.

Before setting a new key identifier, the current key must first be deleted on the Show MD5 Key page.

Web Interface

To configure OSPF interface for all areas assigned to a VLAN:

1. Click Routing Protocol, OSPF, Interface.
2. Select Configure by VLAN from the Action list.
3. Specify the VLAN ID, and configure the required interface settings.
4. Click Apply.

Figure 395: Configuring Settings for All Interfaces Assigned to a VLAN

Routing Protocol > OSPF > Interface

Action: Configure by VLAN

VLAN ID 1

Cost (1-65535)

Router Priority (0-255)

Hello Interval (1-65535) sec

Dead Interval (1-65535) sec

Transmit Delay (1-65535) sec

Retransmit Interval (1-65535) sec

Authentication Type MD5

Message Digest Key ID

Authentication Key

Click the button to clear the configuration of this VLAN.

To configure interface settings for a specific area assigned to a VLAN:

1. Click Routing Protocol, OSPF, Interface.
2. Select Configure by Address from the Action list.
3. Specify the VLAN ID, enter the address assigned to an area, and configure the required interface settings.
4. Click Apply.

Figure 396: Configuring Settings for a Specific Area Assigned to a VLAN

Routing Protocol > OSPF > Interface

Action:

VLAN ID:

IP Address:

Cost (1-65535):

Router Priority (0-255):

Hello Interval (1-65535): sec

Dead Interval (1-65535): sec

Transmit Delay (1-65535): sec

Retransmit Interval (1-65535): sec

Authentication Type:

Message Digest Key ID:

Authentication Key:

Click the button to clear the configuration of this IP address.

To show the configuration settings for OSPF interfaces:

1. Click Routing Protocol, OSPF, Interface.
2. Select Show from the Action list.
3. Select the VLAN ID.

Figure 397: Showing OSPF Interfaces

Routing Protocol > OSPF > Interface

Action: Show

VLAN ID: 1

Interface List Total: 4

Interface IP	Area ID	State	Designated Router IP	Designated Router ID	Backup Designated Router IP	Backup Designated Router ID
192.168.1.2/24	192.168.1.0	Up	192.168.1.2	192.168.1.2	0.0.0.0	0.0.0.0
192.168.10.2/24	192.168.10.0	Up	192.168.10.2	192.168.1.2	0.0.0.0	0.0.0.0
192.168.100.2/24	192.168.100.0	Up	192.168.100.2	192.168.1.2	0.0.0.0	0.0.0.0
192.168.110.2/24	192.168.110.0	Up	192.168.110.2	192.168.1.2	0.0.0.0	0.0.0.0

To show the MD5 authentication keys configured for an interface:

1. Click Routing Protocol, OSPF, Interface.
2. Select Show MD5 Key from the Action list.
3. Select the VLAN ID.

Figure 398: Showing MD5 Authentication Keys

Routing Protocol > OSPF > Interface

Action: Show MD5 Key

VLAN ID: 1

Interface MD5 List Total: 2

<input type="checkbox"/>	Area ID	Key ID
<input type="checkbox"/>	0.0.0.0	1
<input type="checkbox"/>	192.168.10.0	2

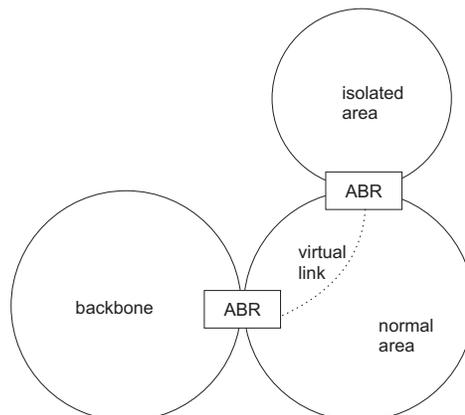
Apply Revert

Configuring Virtual Links

Use the Routing Protocol > OSPF > Virtual Link (Add) and (Configure Detailed Settings) pages to configure a virtual link from an area that does not have a direct physical connection to the OSPF backbone.

All OSPF areas must connect to the backbone. If an area does not have a direct physical connection to the backbone, you can configure a virtual link that provides a logical path to the backbone. To connect an isolated area to the backbone, the logical path can cross a single non-backbone area (i.e., transit area) to reach the backbone. To define this path, you must configure an ABR that serves as an endpoint connecting the isolated area to the common transit area, and specify a neighboring ABR at the other endpoint connecting the common transit area to the backbone itself. (Note that you cannot configure a virtual link that runs through a stub or NSSA.)

Figure 399: OSPF Virtual Link



Virtual links can also be used to create a redundant link between any area and the backbone to help prevent partitioning, or to connect two existing backbone areas into a common backbone.

Any area disconnected from the backbone must include the transit area ID and the router ID for a virtual link neighbor that is adjacent to the backbone.

This router supports up five virtual links.

Command Usage

- ◆ Use the Add page to create a virtual link, and then use the Configure Detailed Settings page to set the protocol timers and authentication settings for the link. The parameters to be configured on the Configure Detailed Settings page are described under [“Configuring OSPF Interfaces” on page 584](#).

Parameters

These parameters are displayed:

- ◆ **Process ID** – Process ID as configured in the Network Area configuration screen (see [page 563](#)).
- ◆ **Transit Area ID** – Identifies the transit area for the virtual link. The area ID must be in the form of an IPv4 address, or also as a four octet unsigned integer ranging from 0-4294967295.
- ◆ **Neighbor ID** – Router ID of the virtual link neighbor. This specifies the Area Border Router (ABR) at the other end of the virtual link. To create a virtual link, it must be configured for an ABR at both ends of the link. One of the ABRs must be next to the isolated area and the transit area at one end of the link, while the other ABR must be next to the transit area and backbone at the other end of the link.

Web Interface

To create a virtual link:

1. Click Routing Protocol, OSPF, Virtual Link.
2. Select Add from the Action list.
3. Specify the process ID, the Area ID, and Neighbor router ID.
4. Click Apply.

Figure 400: Adding a Virtual Link

Routing Protocol > OSPF > Virtual Link

Action: Add

Process ID: 1

Transit Area ID: 192.168.10.0

Neighbor ID: 192.168.10.3

Apply Revert

To show virtual links:

1. Click Routing Protocol, OSPF, Virtual Link.
2. Select Show from the Action list.
3. Select the process ID.

Figure 401: Showing Virtual Links

Routing Protocol > OSPF > Virtual Link

Action: Show

Process ID: 1

Virtual Link List Total: 2

<input type="checkbox"/>	Transit Area ID	Neighbor ID	State	Local Address	Remote Address	Hello Due	Adjacency State
<input type="checkbox"/>	0.0.0.1	10.2.0.0	Down	192.168.1.1	192.168.2.1	Inactive	Full
<input type="checkbox"/>	0.0.0.2	10.3.0.0	Waiting	*	*		Down

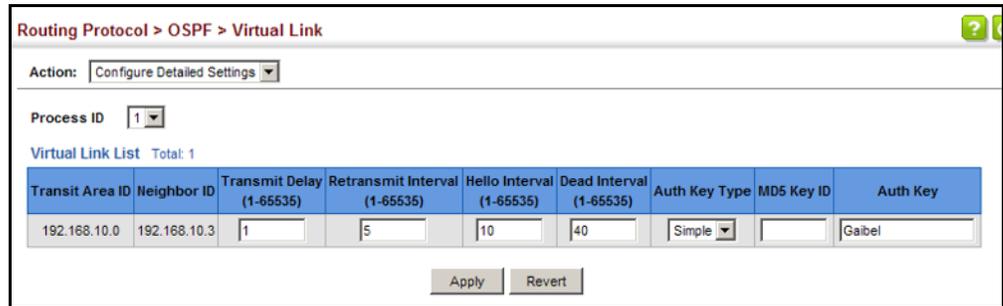
Delete Revert

To configure detailed settings for a virtual link:

1. Click Routing Protocol, OSPF, Virtual Link.
2. Select Configure Detailed Settings from the Action list.
3. Specify the process ID, then modify the protocol timers and authentication settings as required.

4. Click Apply.

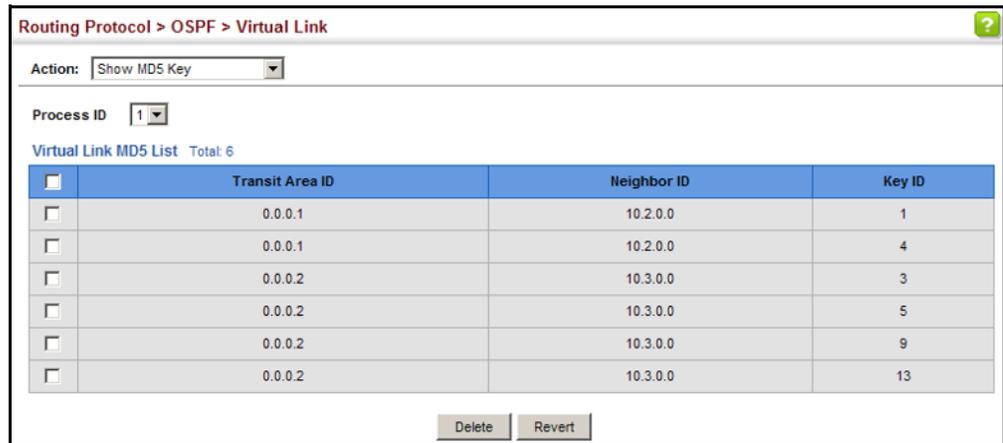
Figure 402: Configuring Detailed Settings for a Virtual Link



To show the MD5 authentication keys configured for a virtual link:

1. Click Routing Protocol, OSPF, Interface.
2. Select Show MD5 Key from the Action list.
3. Select the VLAN ID.

Figure 403: Showing MD5 Authentication Keys



Displaying Link State Database Information

Use the Routing Protocol > OSPF > Information (LSDB) page to show the Link State Advertisements (LSAs) sent by OSPF routers advertising routes. The full collection of LSAs collected by a router interface from the attached area is known as a link state database. Routers that are connected to multiple interfaces will have a separate database for each area. Each router in the same area should have an identical database describing the topology for that area, and the shortest path to external destinations.

The full database is exchanged between neighboring routers as soon as a new router is discovered. Afterwards, any changes that occur in the routing tables are synchronized with neighboring routers through a process called reliable flooding.

You can show information about different LSAs stored in this router's database, which may include any of the following types:

- ◆ Router (Type 1) – All routers in an OSPF area originate Router LSAs that describe the state and cost of its active interfaces and neighbors.
- ◆ Network (Type 2) – The designated router for each area originates a Network LSA that describes all the routers that are attached to this network segment.
- ◆ Summary (Type 3) – Area border routers can generate Summary LSAs that give the cost to a subnetwork located outside the area.
- ◆ AS Summary (Type 4) – Area border routers can generate AS Summary LSAs that give the cost to an autonomous system boundary router (ASBR).
- ◆ AS External (Type 5) – An ASBR can generate an AS External LSA for each known network destination outside the AS.
- ◆ NSSA External (Type 7) – An ASBR within an NSSA generates an NSSA external link state advertisement for each known network destination outside the AS.

Parameters

These parameters are displayed:

- ◆ **Process ID** – Process ID as configured in the Network Area configuration screen (see [page 563](#)).
- ◆ **Query by** – The LSA database can be searched using the following criteria:
 - Self-Originate – LSAs generated by this router.
 - Link ID – LSAs advertising a specific link.
 - Adv Router – LSAs advertised by a specific router.
- ◆ **Link State Type** – The information returned by a query can be displayed for all LSA types or for a specific type. (Default: All)

Information displayed for each LSA entry includes:

- ◆ **Area ID** – Area defined for which LSA information is to be displayed.
- ◆ **Link ID** – Network portion described by an LSA. The Link ID is either:
 - An IP network number for Type 3 Summary and Type 5 AS External LSAs. (When an Type 5 AS External LSA is describing a default route, its Link ID is set to the default destination 0.0.0.0.)
 - A Router ID for Router, Network, and Type 4 AS Summary LSAs.
- ◆ **Adv Router** – IP address of the advertising router.
- ◆ **Age** – Age of LSA (in seconds).

- ◆ **Sequence** – Sequence number of LSA (used to detect older duplicate LSAs).
- ◆ **Checksum** – Checksum of the complete contents of the LSA.

Web Interface

To display information in the link state database:

1. Click Routing Protocol, OSPF, Information.
2. Click LSDB.
3. Select the process identifier.
4. Specify required search criteria, such as self-originated LSAs, LSAs with a specific link ID, or LSAs advertised by a specific router.
5. Then select the database entries to display based on LSA type.

Figure 404: Displaying Information in the Link State Database

The screenshot shows the 'Routing Protocol > OSPF > Information' web interface. The 'Type' is set to 'LSDB' and 'Process ID' is '1'. Under 'Query by', 'Self-Originate', 'Link ID', and 'Adv Router' are all unchecked. The 'Link State Type' is set to 'All'. Below the query options are several tables showing the results of the query.

Link State Router List Total: 2

Area ID	Link ID	Adv Router	Age	Sequence	Checksum
0.0.0.0	192.168.0.4	192.168.0.4	702	0x80000003	0xE6B4
0.0.0.0	192.168.1.2	192.168.1.2	355	0x80000005	0xDDBC

Link State Network List Total: 1

Area ID	Link ID	Adv Router	Age	Sequence	Checksum
0.0.0.0	192.168.0.4	192.168.0.4	702	0x80000001	0x8F16

Link State Summary List Total: 1

Area ID	Link ID	Adv Router	Age	Sequence	Checksum
0.0.0.0	192.168.1.0	192.168.1.2	638	0x80000001	0x99EB

Link State ASBR Summary List Total: 0

Area ID	Link ID	Adv Router	Age	Sequence	Checksum
---------	---------	------------	-----	----------	----------

Link State External List Total: 0

Area ID	Link ID	Adv Router	Age	Sequence	Checksum
---------	---------	------------	-----	----------	----------

Link State NSSA External List Total: 0

Area ID	Link ID	Adv Router	Age	Sequence	Checksum
---------	---------	------------	-----	----------	----------

Displaying Information on Neighboring Routers

Use the Routing Protocol > OSPF > Information (Neighbor) page to display information about neighboring routers on each interface.

Parameters

These parameters are displayed:

- ◆ **Process ID** – Process ID as configured in the Network Area configuration screen (see [page 563](#)).
- ◆ **ID** – Neighbor's router ID.
- ◆ **Priority** – Neighbor's router priority.
- ◆ **State** – OSPF state and identification flag.

States include:

- Down – Connection down
- Attempt – Connection down, but attempting contact (non-broadcast networks)
- Init – Have received Hello packet, but communications not yet established
- Two-way – Bidirectional communications established
- ExStart – Initializing adjacency between neighbors
- Exchange – Database descriptions being exchanged
- Loading – LSA databases being exchanged
- Full – Neighboring routers now fully adjacent

Identification flags include:

- D – Dynamic neighbor
 - S – Static neighbor
 - DR – Designated router
 - BDR – Backup designated router
- ◆ **Address** – IP address of this interface.
 - ◆ **Interface** – A Layer 3 interface on which OSPF has been enabled.

Web Interface

To display information about neighboring routers stored in the link state database:

1. Click Routing Protocol, OSPF, Information.
2. Click Neighbor.
3. Select the process identifier.

Figure 405: Displaying Neighbor Routers Stored in the Link State Database

The screenshot shows the 'Routing Protocol > OSPF > Information' page. At the top, there are radio buttons for 'LSDB' (unselected) and 'Neighbor' (selected). Below that is a 'Process ID' dropdown menu set to '1'. The main content is a table titled 'Neighbor Information List' with a 'Total: 2' indicator. The table has five columns: Neighbor ID, Priority, State, Address, and Interface. It contains two rows of data.

Neighbor ID	Priority	State	Address	Interface
10.10.10.50	1	Full / DR	10.10.10.50	VLAN 1
10.10.10.50	1	Full / Backup	10.10.10.50	VLAN 2

Specifying Passive Interfaces

Use the Routing Protocol > OSPF > Passive Interface (Add) page to stop OSPF from sending routing updates on the specified interface.

Command Usage

You can configure an OSPF interface as passive to prevent OSPF routing traffic from exiting or entering that interface. No OSPF adjacency can be formed if one of the interfaces involved is set to passive mode. The specified interface will appear as a stub in the OSPF domain. Also, if you configure an OSPF interface as passive where an adjacency already exists, the adjacency will drop almost immediately.

Parameters

These parameters are displayed:

- ◆ **Process ID** – Protocol identifier as configured on the Routing Protocol > OSPF > Network Area (Add) page. (Range: 1-65535)
- ◆ **VLAN ID** – VLAN ID. (Range: 1-4094)
- ◆ **IP Address** – An IPv4 address configured on this interface.

Web Interface

To specify a passive OSPF interface:

1. Click Routing Protocol, OSPF, Passive Interface.
2. Select Add from the Action list.
3. Add the interface on which to stop sending OSPF routing traffic..
4. Click Apply.

Figure 406: Specifying a Passive OSPF Interface

Routing Protocol > OSPF > Passive Interface

Action: Add ▾

Process ID 1 ▾

VLAN ID 1 ▾

IP Address 192.168.0.4

Apply Revert

To show the passive OSPF interfaces:

1. Click Routing Protocol, OPPE, Passive Interface.
2. Select Show from the Action list.

Figure 407: Showing Passive OSPF Interfaces

Routing Protocol > OSPF > Passive Interface

Action: Show ▾

OSPF Passive Interface List Total: 1

<input type="checkbox"/>	Process ID	VLAN	IP Address
<input type="checkbox"/>	1	1	192.168.0.4

Delete Revert

Multicast Routing

This chapter describes the following multicast routing topics:

- ◆ [Enabling Multicast Routing Globally](#) – Describes how to globally enable multicast routing.
- ◆ [Displaying the Multicast Routing Table](#) – Describes how to display the multicast routing table.
- ◆ [Configuring PIM for IPv4](#) – Describes how to configure PIM-DM and PIM-SM for IPv4.
- ◆ [Configuring PIMv6 for IPv6](#) – Describes how to configure PIM-DM and PIM-SM (Version 6) for IPv6.

Overview

This router can route multicast traffic to different subnetworks using Protocol-Independent Multicasting - Dense Mode or Sparse Mode (PIM-DM or PIM-SM) for IPv4, as well as PIM-DM for IPv6. PIM for IPv4 (also called PIMv4 in this manual) relies on messages sent from IGMP-enabled Layer 2 switches and hosts to determine when hosts want to join or leave multicast groups. PIM for IPv6 (also called PIMv6 in this manual) uses the Multicast Listener Discovery (MLDv1) protocol which is the IPv6 equivalent to IGMPv2. PIM-DM is designed for networks where the probability of multicast group members is high, such as a local network. PIM-SM is designed for networks where the probability of multicast group members is low, such as the Internet.

Also, note that if PIM is not enabled on this router or another multicast routing protocol is used on the network, the switch ports attached to a multicast router can be manually configured to forward multicast traffic (see [“Specifying Static Interfaces for an IPv4 Multicast Router” on page 441](#)).

Configuring PIM-DM

PIM-DM floods multicast traffic downstream, and calculates the shortest-path, source-rooted delivery tree between each source and destination host group. Other multicast routing protocols, such as DVMRP, build their own source-rooted multicast delivery tree (i.e., a separate routing table) that allows it to prevent looping and determine the shortest path to the source of the multicast traffic. PIM-DM also builds a source-rooted multicast delivery tree for each multicast source,

but uses information from the router's unicast routing table, instead of maintaining its own multicast routing table, making it routing protocol independent.

PIM-DM is a simple multicast routing protocol that uses flood and prune to build a source-routed multicast delivery tree for each multicast source-group pair. As mentioned above, it does not maintain its own routing table, but instead, uses the routing table provided by whatever unicast routing protocol is enabled on the router interface. When the router receives a multicast packet for a source-group pair, PIM-DM checks the unicast routing table on the inbound interface to determine if this is the same interface used for routing unicast packets to the multicast source network. If it is not, the router drops the packet and sends an Assert message back out the source interface. An Assert winner is then selected to continue forwarding traffic from this source. On the other hand, if it is the same interface used by the unicast protocol, then the router forwards a copy of the packet to all the other interfaces for which it has not already received a prune message for this specific source-group pair.

DVMRP holds the prune state for about two hours, while PIM-DM holds it for only about three minutes. Although this results in more flooding than encountered with DVMRP, this is the only major trade-off for the lower processing overhead and simplicity of configuration for PIM-DM.

Configuring PIM-SM

PIM-SM uses the router's local unicast routing table to route multicast traffic, not to flood it. It only forwards multicast traffic when requested by a local or downstream host. When service is requested by a host, it can use a Reverse Path Tree (RPT) that channels the multicast traffic from each source through a single Rendezvous Point (RP) within the local PIM-SM domain, and then forwards this traffic to the Designated Router (DR) in the local network segment to which the host is attached. However, when the multicast load from a particular source is heavy enough to justify it, PIM-SM can be configured to construct a Shortest Path Tree (SPT) directly from the DR up to the source, bypassing the RP and thereby reducing service delays for active hosts and setup time for new hosts.

PIM-SM reduces the amount of multicast traffic by forwarding it only to the ports that are attached to receivers for a group. The key components to filtering multicast traffic are listed below.

Common Domain – A common domain must be set up in which all of the multicast routers are configured with the same basic PIM-SM settings.

Bootstrap Router (BSR) – After the common domain is set, a bootstrap router is elected from this domain. Each time a PIM-SM router is booted up, or the multicast mode reconfigured to enable PIM-SM, the bootstrap router candidates start flooding bootstrap messages on all of their interfaces (using reverse path forwarding to limit the impact on the network). When neighboring routers receive bootstrap messages, they process the message and forward it out through all interfaces, except for the interface on which this message was received. If a router receives a bootstrap message with a BSR priority larger than its own, it stops

advertising itself as a BSR candidate. Eventually, only the router with the highest BSR priority will continue sending bootstrap messages.

Rendezvous Point (RP) – A router may periodically send PIMv2 messages to the BSR advertising itself as a candidate RP for specified group addresses. The BSR places information about all of the candidate RPs in subsequent bootstrap messages. The BSR and all the routers receiving these messages use the same hash algorithm to elect an RP for each multicast group. If each router is properly configured, the results of the election process will be the same for each router. Each elected RP then starts to serve as the root of a shared distribution tree for one or more multicast groups.

Designated Router (DR) – A DR advertising the highest priority in its hello messages is elected for each subnet. The DR is responsible for collecting information from the subnet about multicast clients that want to join or leave a group. Join messages from the DR (receiver) for each group are sent towards the RP, and data from multicast sources is sent to the RP. Receivers can now start receiving traffic destined for the client group from the RP, or they can identify the senders and optionally set up a direct connection to the source through a shortest path tree (SPT) if the loading warrants this change over.

Shared Tree – When many receivers join a group, their Join messages converge on the RP, and form a distribution tree for the group that is rooted at the RP. This is known as the Reverse Path Tree (RPT), or the shared tree since it is shared by all sources sending to that group. When a multicast source sends data destined for a group, the source's local DR takes those data packets, unicast-encapsulates them, and sends them to the RP. When the RP receives these encapsulated data packets, it decapsulates them, and forwards them onto the shared tree. These packets follow the group mapping maintained by routers along the RP Tree, are replicated wherever the RP Tree branches, and eventually reach all the receivers for that multicast group. Because all routers along the shared tree are using PIM-SM, the multicast flow is confined to the shared tree. Also, note that more than one flow can be carried over the same shared tree, but only one RP is responsible for each flow.

Shortest Path Tree (SPT) – When using the Shared Tree, multicast traffic is contained within the shared tree. However, there are several drawbacks to using the shared tree. Decapsulation of traffic at the RP into multicast packets is a resource intensive process. The protocol does not take into account the location of group members when selecting the RP, and the path from the RP to the receiver is not always optimal. Moreover, a high degree of latency may occur for hosts wanting to join a group because the RP must wait for a register message from the DR before setting up the shared tree and establishing a path back to the source. There is also a problem with bursty sources. When a source frequently times out, the shared tree has to be rebuilt each time, causing further latency in sending traffic to the receiver. To enhance overall network performance, the switch uses the RP only to forward the first packet from a source to the receivers. After the first packet, it calculates the shortest path between the receiver and source and uses the SPT to send all subsequent packets from the source directly to the receiver. When the first packet arrives natively through the shortest path, the RP sends a register-stop message back to the DR near the source. When this DR receives the

register-stop message, it stops sending register messages to the RP. If there are no other sources using the shared tree, it is also torn down. Setting up the SPT requires more memory than when using the shared tree, but can significantly reduce group join and data transmission delays. The switch can also be configured to use SPT only for specific multicast groups, or to disable the change over to SPT for specific groups.

Configuring Global Settings for Multicast Routing

To use multicast routing on this router, first globally enable multicast routing as described in this section, then specify the interfaces that will employ multicast routing protocols (PIM-DM or PIM-SM). Note that only one multicast routing protocol (PIM-DM or PIM-SM) can be enabled on any given interface, but both PIMv4 and PIMv6 can be enabled on the same interface.

Enabling Multicast Routing Globally

Use the Multicast > Multicast Routing > General page or the Multicast > IPv6 Multicast Routing > General page to enable IPv4 or IPv6 multicast routing globally on the switch.

Parameters

These parameters are displayed:

IPv4 Multicast Routing

- ◆ **Multicast Forwarding Status** – Enables IP multicast routing. (Default: Disabled)

IPv6 Multicast Routing

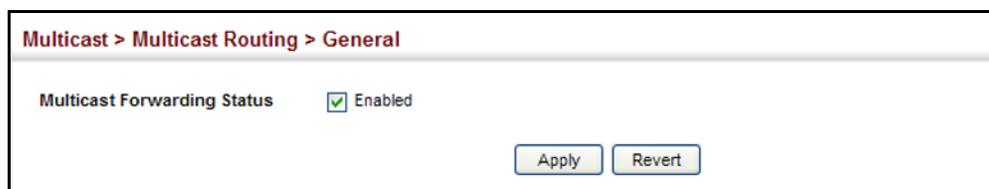
- ◆ **IPv6 Multicast Forwarding Status** – Enables IPv6 multicast routing. (Default: Disabled)

Web Interface (IPv4)

To enable IPv4 multicast routing:

1. Click Multicast, Multicast Routing, General.
2. Enable Multicast Forwarding Status.
3. Click Apply.

Figure 408: Enabling IPv4 Multicast Routing

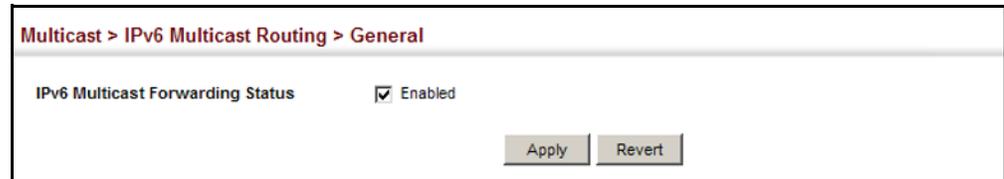


Web Interface (IPv6)

To enable IPv6 multicast routing:

1. Click Multicast, IPv6 Multicast Routing, General.
2. Enable Multicast Forwarding Status.
3. Click Apply.

Figure 409: Enabling IPv6 Multicast Routing



Displaying the Multicast Routing Table

Use the Multicast > Multicast Routing > Information page or the IPv6 Multicast > Multicast Routing > Information page to display IPv4 or IPv6 information on each multicast route the switch has learned through PIM. The router learns multicast routes from neighboring routers, and also advertises these routes to its neighbors. The router stores entries for all paths learned by itself or from other routers, without considering actual group membership or prune messages. The routing table therefore does not indicate that the router has processed multicast traffic from any particular source listed in the table. It uses these routes to forward multicast traffic only if group members appear on directly-attached subnetworks or on subnetworks attached to downstream routers.

Parameters

These parameters are displayed for IPv4:

Show Summary

- ◆ **Group Address** – IP group address for a multicast service.
- ◆ **Source Address** – Subnetwork containing the IP multicast source.
- ◆ **Source Mask** – Network mask for the IP multicast source. Note that the switch cannot detect the source mask, and therefore displays 255.255.255.255 in this field. (This parameter applies to IPv4 only.)
- ◆ **Interface** – Upstream interface leading to the upstream neighbor.
PIM creates a multicast routing tree based on the unicast routing table. If the related unicast routing table does not exist, PIM will still create a multicast routing entry, displaying the upstream interface to indicate that this entry is valid. This field may also display “Register” to indicate that a pseudo interface is being used to receive PIM-SM register packets. This can occur for the Rendezvous Point (RP), which is the root of the Reverse Path Tree (RPT). In this

case, any VLAN receiving register packets will be converted into the register interface.

- ◆ **Owner** – The associated multicast protocol (PIM-DM, PIM-SM, IGMP Proxy for PIMv4, MLD Proxy for PIMv6).
- ◆ **Flags** – The flags associated with each routing entry indicate:
 - **Forward** – Traffic received from the upstream interface is being forwarded to this interface.
 - **Local** – This is the outgoing interface.
 - **Pruned** – This interface has been pruned by a downstream neighbor which no longer wants to receive the traffic.

Show Details

- ◆ **Group Address** – IP group address for a multicast service.
- ◆ **Source Address** – Subnetwork containing the IP multicast source.
- ◆ **Source Mask** – Network mask for the IP multicast source.
- ◆ **Upstream Neighbor** – The multicast router (RPF Neighbor) immediately upstream for this group.
- ◆ **Upstream Interface** – Interface leading to the upstream neighbor.
- ◆ **Up Time** – Time since this entry was created.
- ◆ **Owner** – The associated multicast protocol (PIM-DM, PIM-SM, IGMP Proxy for PIMv4, MLD Proxy for PIMv6).
- ◆ **Flags** – The flags associated with each routing entry indicate:
 - **Dense** – PIM Dense mode in use.
 - **Sparse** – PIM Sparse mode in use.
 - **Connected** – This route is directly connected to the source.
 - **Pruned** – This route has been terminated.
 - **Register flag** – This device is registering for a multicast source.
 - **RPT-bit set** – The (S,G) entry is pointing to the Rendezvous Point (RP), which normally indicates a pruned state along the shared tree for a particular source.

- **SPT-bit set** – Multicast packets have been received from a source on shortest path tree.
- **Join SPT** – The rate of traffic arriving over the shared tree has exceeded the SPT-threshold for this group. If the SPT flag is set for (*,G) entries, the next (S,G) packet received will cause the router to join the shortest path tree. If the SPT flag is set for (S,G), the router immediately joins the shortest path tree.

Downstream Interface List

- ◆ **Interface** – Interface(s) on which multicast subscribers have been recorded.
- ◆ **State** – The flags associated with each downstream interface indicate:
 - **Forward** – Traffic received from the upstream interface is being forwarded to this interface.
 - **Local** – Downstream interface has received IGMP report message from host in this subnet.
 - **Pruned** – This route has been terminated.
 - **Registering** – A downstream device is registering for a multicast source.

Web Interface (IPv4)

To display the multicast routing table:

1. Click Multicast, Multicast Routing, Information.
2. Select Show Summary from the Action List.

Figure 410: Displaying the IPv4 Multicast Routing Table

Multicast > Multicast Routing > Information					
Action: <input type="button" value="Show Summary"/>					
Multicast Routing Summary List Total: 3					
Group Address	Source Address	Source Mask	Interface	Owner	Flags
224.0.17.17	192.168.2.1	255.255.255.255	VLAN 1	PIM-DM	Forward
224.1.1.1	10.1.1.0	255.255.255.0	VLAN 2	DVMRP	Pruned
224.1.1.2	10.1.1.0	255.255.255.0	VLAN 3	DVMRP	Forward

To display detailed information on a specific flow in multicast routing table:

1. Click Multicast, Multicast Routing, Information.
2. Select Show Details from the Action List.
3. Select a Group Address.

4. Select a Source Address.

Figure 411: Displaying Detailed Entries from IPv4 Multicast Routing Table

The screenshot shows the configuration page for Multicast Routing. The breadcrumb is "Multicast > Multicast Routing > Information". The "Action" dropdown is set to "Show Details". The configuration parameters are as follows:

Group Address	224.0.17.17
Source Address	192.168.2.1
Source Mask	255.255.255.255
Upstream Neighbor	192.168.2.2
Upstream Interface	VLAN 1
Up Time	00:00:05
Owner	PIM-DM
Flags	Dense

Below the configuration parameters is a "Downstream Interface List" with a total of 3 entries:

Interface	State
VLAN 1	Forward
VLAN 2	Pruned
VLAN 3	Forward

Web Interface (IPv6)

To display the multicast routing table:

1. Click Multicast, IPv6 Multicast Routing, Information.
2. Select Show Summary from the Action List.

Figure 412: Displaying the IPv6 Multicast Routing Table

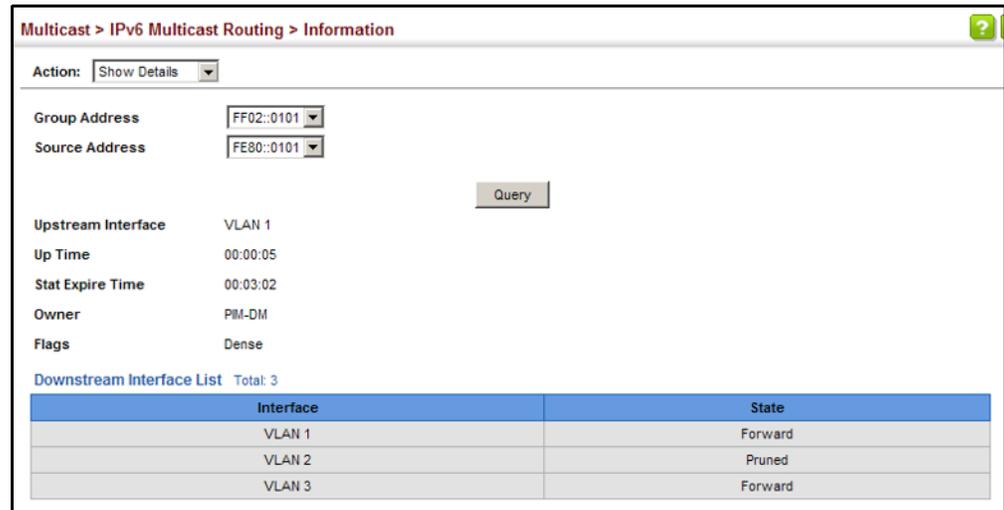
The screenshot shows the configuration page for IPv6 Multicast Routing. The breadcrumb is "Multicast > IPv6 Multicast Routing > Information". The "Action" dropdown is set to "Show Summary". The "IPv6 Multicast Routing Summary List" has a total of 3 entries:

Group Address	Source Address	Interface	Owner	Flags
FF02::0101	FE80::0101	VLAN 4096	PIM-DM	Forward
FF02::0102	FE80::0102	VLAN 4095	PIM-DM	Forward
FF02::0103	FE80::0103	VLAN 4094	PIM-SM	Pruned

To display detailed information on a specific flow in multicast routing table:

1. Click Multicast, IPv6 Multicast Routing, Information.
2. Select Show Details from the Action List.
3. Select a Group Address.
4. Select a Source Address.
5. Click Query.

Figure 413: Displaying Detailed Entries from IPv6 Multicast Routing Table



Configuring PIM for IPv4

This section describes how to configure PIM-DM and PIM-SM for IPv4.

Enabling PIM Globally Use the Routing Protocol > PIM > General page to enable IPv4 PIM routing globally on the router.

Command Usage

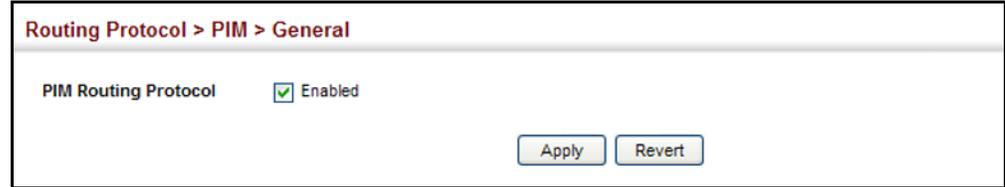
- ◆ This feature enables PIM-DM and PIM-SM globally for the router. You also need to enable PIM-DM or PIM-SM for each interface that will support multicast routing (see [page 608](#)), and make any changes necessary to the multicast protocol parameters.
- ◆ To use PIM, multicast routing must be enabled on the switch (see ["Enabling Multicast Routing Globally" on page 602](#)).

Web Interface

To enable PIM multicast routing:

1. Click Routing Protocol, PIM, General.
2. Enable PIM Routing Protocol.
3. Click Apply.

Figure 414: Enabling PIM Multicast Routing



Configuring PIM Interface Settings Use the Routing Protocol > PIM > Interface page configure the routing protocol's functional attributes for each interface.

Command Usage

- ◆ Most of the attributes on this page are common to both PIM-DM and PIM-SM. Select Dense or Sparse Mode to display the common attributes, as well as those applicable to the selected mode.
- ◆ PIM and IGMP proxy cannot be used at the same time. When an interface is set to use PIM Dense mode or Sparse mode, IGMP proxy cannot be enabled on any interface of the device (see [“Configuring IGMP Snooping and Query Parameters” on page 437](#)). Also, when IGMP proxy is enabled on an interface, PIM cannot be enabled on any interface.

PIM-DM

- ◆ PIM-DM functions similar to DVMRP by periodically flooding the network with traffic from any active multicast server. It also uses IGMP to determine the presence of multicast group members. The main difference, is that it uses the router's unicast routing table to determine if the interface through which a packet is received provides the shortest path back to the source.
- ◆ Dense-mode interfaces are subject to multicast flooding by default, and are only removed from the multicast routing table when the router determines that there are no group members or downstream routers, or when a prune message is received from a downstream router.

PIM-SM

- ◆ A PIM-SM interface is used to forward multicast traffic only if a join message is received from a downstream router or if group members are directly connected to the interface. When routers want to receive a multicast flow, they periodically send join messages to the RP, and are subsequently added to the shared path for the specified flow back up to the RP. If routers want to join the source path up through the SPT, they periodically send join messages toward the source. They also send prune messages toward the RP to prune the shared path once they have connected to the source through the SPT, or if there are no longer any group members connected to the interface.

Parameters

These parameters are displayed:

Common Attributes

- ◆ **VLAN** – Layer 3 VLAN interface. (Range: 1-4094)
- ◆ **Mode** – PIM routing mode. (Options: Dense, Sparse, None)
- ◆ **IP Address** – Primary IP address assigned to the selected VLAN.
- ◆ **Hello Holdtime** – Sets the interval to wait for hello messages from a neighboring PIM router before declaring it dead. Note that the hello holdtime should be greater than or equal to the value of Hello Interval, otherwise it will be automatically set to 3.5 x the Hello Interval. (Range: 1-65535 seconds; Default: 105 seconds, or 3.5 times the hello interval if set)
- ◆ **Hello Interval** – Sets the frequency at which PIM hello messages are transmitted out on all interfaces. (Range: 1-65535 seconds; Default: 30 seconds)

Hello messages are sent to neighboring PIM routers from which this device has received probes, and are used to verify whether or not these neighbors are still active members of the multicast tree. PIM-SM routers use these messages not only to inform neighboring routers of their presence, but also to determine which router for each LAN segment will serve as the Designated Router (DR).

When a router is booted or first configured to use PIM, it sends an initial hello message, and then sets its Hello timer to the configured value. If a router does not hear from a neighbor for the period specified by the Hello Holdtime, that neighbor is dropped. This hold time is included in each hello message received from a neighbor. Also note that hello messages also contain the DR priority of the router sending the message.

If the hello holdtime is already configured, and the hello interval is set to a value longer than the hello holdtime, this command will fail.

- ◆ **Join/Prune Holdtime** – Sets the hold time for the prune state. (Range: 1-65535 seconds; Default: 210 seconds)
 - PIM-DM: The multicast interface that first receives a multicast stream from a particular source forwards this traffic to all other PIM-DM interfaces on the router. If there are no requesting groups on that interface, the leaf node sends a prune message upstream and enters a prune state for this multicast stream. The prune state is maintained until the join/prune holdtime timer expires or a graft message is received for the forwarding entry.
 - PIM-SM: The multicast interface that first receives a multicast stream from a particular source forwards this traffic only to those interfaces on the router that have requests to join this group. When there are no longer any requesting groups on that interface, the leaf node sends a prune message upstream and enters a prune state for this multicast stream. The protocol maintains both the current join state and the pending RPT prune state for this (source, group) pair until the join/prune interval timer expires.

- ◆ **LAN Prune Delay** – Causes this device to inform downstream routers of how long it will wait before pruning a flow after receiving a prune request. (Default: Disabled)

When other downstream routers on the same VLAN are notified that this upstream router has received a prune request, they must send a Join to override the prune before the prune delay expires if they want to continue receiving the flow. The message generated by this command effectively prompts any downstream neighbors with hosts receiving the flow to reply with a Join message. If no join messages are received after the prune delay expires, this router will prune the flow.

The sum of the Override Interval and Propagation Delay are used to calculate the LAN prune delay.

- ◆ **Override Interval** – The time required for a downstream router to respond to a LAN Prune Delay message by sending back a Join message if it wants to continue receiving the flow referenced in the message. (Range: 500-6000 milliseconds; Default: 2500 milliseconds)

The override interval and the propagation delay are used to calculate the LAN prune delay. If a downstream router has group members which want to continue receiving the flow referenced in a LAN prune delay message, then the override interval represents the time required for the downstream router to process the message and then respond by sending a Join message back to the upstream router to ensure that the flow is not terminated.

- ◆ **Propagation Delay** – The time required for a LAN prune delay message to reach downstream routers. (Range: 100-5000 milliseconds; Default: 500 milliseconds)

The override interval and propagation delay are used to calculate the LAN prune delay. If a downstream router has group members which want to continue receiving the flow referenced in a LAN prune delay message, then the propagation delay represents the time required for the LAN prune delay message to be propagated down from the upstream router to all downstream routers attached to the same VLAN interface.

- ◆ **Trigger Hello Delay** – The maximum time before transmitting a triggered PIM Hello message after the router is rebooted or PIM is enabled on an interface. (Range: 0-5 seconds; Default: 5 seconds)

When a router first starts or PIM is enabled on an interface, the hello delay is set to random value between 0 and the trigger hello delay. This prevents synchronization of Hello messages on multi-access links if multiple routers are powered on simultaneously.

Also, if a Hello message is received from a new neighbor, the receiving router will send its own Hello message after a random delay between 0 and the trigger hello delay.

Dense-Mode Attributes

- ◆ **Graft Retry Interval** – The time to wait for a Graft acknowledgement before resending a Graft message. (Range: 1-10 seconds; Default: 3 seconds)

A graft message is sent by a router to cancel a prune state. When a router receives a graft message, it must respond with an graft acknowledgement message. If this acknowledgement message is lost, the router that sent the graft message will resend it a number of times (as defined by Max. Graft Retries).

- ◆ **Max. Graft Retries** – The maximum number of times to resend a Graft message if it has not been acknowledged. (Range: 1-10; Default: 3)

- ◆ **State Refresh Origination Interval** – The interval between sending PIM-DM state refresh control messages. (Range: 1-100 seconds; Default: 60 seconds)

The pruned state times out approximately every three minutes and the entire PIM-DM network is re-flooded with multicast packets and prune messages. The state refresh feature keeps the pruned state from timing out by periodically forwarding a control message down the distribution tree, refreshing the prune state on the outgoing interfaces of each router in the tree. This also enables PIM routers to recognize topology changes (sources joining or leaving a multicast group) before the default three-minute state timeout expires.

This command is only effectively for interfaces of first hop, PIM-DM routers that are directly connected to the sources of multicast groups.

Sparse-Mode Attributes

- ◆ **DR Priority** – Sets the priority advertised by a router when bidding to become the Designated Router (DR). (Range: 0-4294967294; Default: 1)

More than one PIM-SM router may be connected to an Ethernet or other shared-media LAN. If multicast hosts are directly connected to the LAN, then only one of these routers is elected as the DR, and acts on behalf of these hosts, sending periodic Join/Prune messages toward a group-specific RP for each group. A single DR is elected per interface (LAN or otherwise) using a simple election process.

The router with the highest priority configured on an interface is elected as the DR. If more than one router attached to this interface uses the same priority, then the router with the highest IP address is elected to serve as the DR.

If a router does not advertise a priority in its hello messages, it is assumed to have the highest priority and is elected as the DR. If more than one router is not advertising its priority, then the router with the highest IP address is elected to serve as the DR.

- ◆ **Join/Prune Interval** – Sets the interval at which join/prune messages are sent. (Range: 1-65535 seconds; Default: 60 seconds)

By default, the switch sends join/prune messages every 60 seconds to inform other PIM-SM routers about clients who want to join or leave a multicast group.

Use the same join/prune message interval on all PIM-SM routers in the same PIM-SM domain, otherwise the routing protocol's performance will be adversely affected.

The multicast interface that first receives a multicast stream from a particular source forwards this traffic only to those interfaces on the router that have requests to join this group. When there are no longer any requesting groups on that interface, the leaf node sends a prune message upstream and enters a prune state for this multicast stream. The protocol maintains both the current join state and the pending RPT prune state for this (source, group) pair until the join/prune interval timer expires.

Web Interface

To configure PIM interface settings:

1. Click Routing Protocol, PIM, Interface.
2. Modify any of the protocol parameters as required.
3. Click Apply.

Figure 415: Configuring PIM Interface Settings (Dense Mode)

The screenshot shows a web interface titled "Routing Protocol > PIM > Interface". The interface contains the following configuration fields:

VLAN	1
Mode	Dense
IP Address	192.168.0.2
Hello Holdtime (1-65535)	105 sec
Hello Interval (1-65535)	30 sec
Join/Prune Holdtime (1-65535)	210 sec
LAN Prune Delay	<input type="checkbox"/> Enabled
Override Interval (500-6000)	2500 msec
Propagation Delay (100-5000)	500 msec
Trigger Hello Delay (0-5)	5 sec
Graft Retry Interval (1-10)	3 sec
Max. Graft Retries (1-10)	3
State Refresh Origination Interval (1-100)	60 sec

At the bottom right of the form, there are two buttons: "Apply" and "Revert".

Figure 416: Configuring PIM Interface Settings (Sparse Mode)

Routing Protocol > PIM > Interface

VLAN: 1

Mode: Sparse

IP Address: 192.168.0.2

Hello Holdtime (1-65535): 105 sec

Hello Interval (1-65535): 30 sec

Join/Prune Holdtime (1-65535): 210 sec

LAN Prune Delay: Enabled

Override Interval (500-6000): 2500 msec

Propagation Delay (100-5000): 500 msec

Trigger Hello Delay (0-5): 5 sec

DR Priority (0-4294967294): 1

Join/Prune Interval (1-65535): 60 sec

Apply Revert

Displaying PIM Neighbor Information

Use the Routing Protocol > PIM > Neighbor page to display all neighboring PIM routers.

Parameters

These parameters are displayed:

- ◆ **Address** – IP address of the next-hop router.
- ◆ **VLAN** – VLAN that is attached to this neighbor.
- ◆ **Uptime** – The duration this entry has been active.
- ◆ **Expire** – The time before this entry will be removed.
- ◆ **DR** – Indicates if a neighbor is the designated router.

Web Interface

To display neighboring PIM routers:

1. Click Routing Protocol, PIM, Neighbor.

Figure 417: Showing PIM Neighbors

Routing Protocol > PIM > Neighbor

Neighbor Information Total: 2

Address	VLAN	Uptime	Expire	DR
10.1.2.50	1	00:01:23	00:01:23	Yes
10.1.2.51	2	1d11h	Never	Yes

Configuring Global PIM-SM Settings

Use the Routing Protocol > PIM > PIM-SM (Configure Global) page to configure the rate at which register messages are sent, the source of register messages, and switch over to the Shortest Path Tree (SPT).

Parameters

These parameters are displayed:

- ◆ **Register Rate Limit** – Configures the rate at which register messages are sent by the Designated Router (DR) for each (source, group) entry. (Range: 1-65535 packets per second; Default: disabled)

This parameter can be used to relieve the load on the designated router (DR) and rendezvous point (RP). However, because register messages exceeding the limit are dropped, some receivers may experience data packet loss within the first few seconds in which register messages are sent from bursty sources.

- ◆ **Register Source** – Configures the IP source address of a register message to an address other than the outgoing interface address of the DR that leads back toward the RP. (Range: VLAN 1-4094; Default: The IP address of the DR's outgoing interface that leads back to the RP)

When the source address of a register message is filtered by intermediate network devices, or is not a uniquely routed address to which the RP can send packets, the replies sent from the RP to the source address will fail to reach the DR, resulting in PIM-SM protocol failures. This type of problem can be overcome by manually configuring the source address of register messages to an interface that leads back to the RP.

- ◆ **SPT Threshold** – Prevents the last-hop PIM-SM router from switching to Shortest Path Source Tree (SPT) mode. (Options: Infinity, Reset; Default: Reset, to use the SPT)

The default path for packets from a multicast source to a receiver is through the RP. However, the path through the RP is not always the shortest path. Therefore, the router uses the RP to forward only the first packet from a new multicast group to its receivers. Afterwards, it calculates the shortest path tree (SPT) directly between the receiver and source, and then uses the SPT to send all subsequent packets from the source to the receiver instead of using the shared tree. Note that when the SPT threshold is not set by this command, the PIM leaf router will join the shortest path tree immediately after receiving the first packet from a new source.

Enable the SPT threshold by selecting "Reset" to force the router to use the shared tree for all multicast groups, or just for the specified multicast groups. (This is the default setting.)

- ◆ **Group Address** – An IP multicast group address. If a group address is not specified, the shared tree is used for all multicast groups.
- ◆ **Group Mask** – Subnet mask that is used for the group address.

Web Interface

To configure global settings for PIM-SM:

1. Click Multicast, Multicast Routing, SM.
2. Select Configure Global from the Step list.
3. Set the register rate limit and source of register messages if required. Also specify any multicast groups which must be routed across the shared tree, instead of switching over to the SPT.
4. Click Apply.

Figure 418: Configuring Global Settings for PIM-SM

Routing Protocol > PIM > SM

Step: 1. Configure Global

Register Rate Limit (1-65535) Enabled 500 packets/sec

Register Source Enabled VLAN 1

SPT Threshold Infinity

Group Address 224.1.0.0 (Optional)

Group Mask 255.255.0.0 (Optional)

Apply Revert

Configuring a PIM BSR Candidate Use the Routing Protocol > PIM > PIM-SM (BSR Candidate) page to configure the switch as a Bootstrap Router (BSR) candidate.

Command Usage

- ◆ When this router is configured as a BSR candidate, it starts sending bootstrap messages to all of its PIM-SM neighbors. The primary IP address of the designated VLAN is sent as the candidate's BSR address. Each neighbor receiving the bootstrap message compares the BSR address with the address from previous messages. If the current address is the same or a higher address, it accepts the bootstrap message and forwards it. Otherwise, it drops the message.
- ◆ This router will continue to be the BSR until it receives a bootstrap message from another candidate with a higher priority (or a higher IP address if the priorities are the same).
- ◆ To improve failover recovery, it is advisable to select at least two core routers in diverse locations, each to serve as both a candidate BSR and candidate RP. It is also preferable to set up one of these routers as both the primary BSR and RP.

Parameters

These parameters are displayed:

- ◆ **BSR Candidate Status** – Configures the switch as a Bootstrap Router (BSR) candidate. (Default: Disabled)
- ◆ **VLAN ID** – Identifier of configured VLAN interface. (Range: 1-4094)
- ◆ **Hash Mask Length** – Hash mask length (in bits) used for RP selection (see “Configuring a PIM Static Rendezvous Point” on page 617 and “Configuring a PIM RP Candidate” on page 618). The portion of the hash specified by the mask length is ANDed with the group address. Therefore, when the hash function is executed on any BSR, all groups with the same seed hash will be mapped to the same RP. If the mask length is less than 32, then only the first portion of the hash is used, and a single RP will be defined for multiple groups. (Range: 0-32; Default: 10)
- ◆ **Priority** – Priority used by the candidate bootstrap router in the election process. The BSR candidate with the largest priority is preferred. If the priority values are the same, the candidate with the larger IP address is elected to be the BSR. Setting the priority to zero means that this router is not eligible to server as the BSR. At least one router in the PIM-SM domain must be set to a value greater than zero. (Range: 0-255; Default: 0)

Web Interface

To configure the switch as a BSR candidate:

1. Click Routing Protocol, PIM, SM.
2. Select BSR Candidate from the Step list.
3. Specify the VLAN interface for which this router is bidding to become the BSR, the hash mask length that will subsequently be used for RP selection if this router is selected as the BSR, and the priority for BSR selection.
4. Click Apply.

Figure 419: Configuring a PIM-SM BSR Candidate

Routing Protocol > PIM > SM

Step: 2. BSR Candidate

BSR Candidate Status Enabled

VLAN ID 1

Hash Mask Length (0-32) 20

Priority (0-255) 200

Apply Revert

Configuring a PIM Static Rendezvous Point

Use the Routing Protocol > PIM > PIM-SM (RP Address) page to configure a static address as the Rendezvous Point (RP) for a particular multicast group.

Command Usage

- ◆ The router will act as an RP for all multicast groups in the local PIM-SM domain if no groups are specified. A static RP can either be configured for the whole multicast group range 224/4, or for specific group ranges.
- ◆ If an IP address is specified that was previously used for an RP, then the older entry is replaced.
- ◆ Multiple RPs can be defined for different groups or group ranges. If a group is matched by more than one entry, the router will use the RP associated with the longer group prefix length. If the prefix lengths are the same, then the static RP with the highest IP address is chosen.
- ◆ Static definitions for RP addresses may be used together with RP addresses dynamically learned through the bootstrap router (BSR). If an RP address learned by the BSR and one statically configured using this command are both available for a group range, the RP address learned by the BSR is chosen over the one statically configured.
- ◆ All routers within the same PIM-SM domain must be configured with the same RP(s). Selecting an RP through the dynamic election process is therefore preferable for most situations. Using the dynamic RP election process also allows a backup RP to automatically take over if the active RP router becomes unavailable.

Parameters

These parameters are displayed:

- ◆ **RP Address** – Static IP address of the router that will be an RP for the specified multicast group(s).
- ◆ **Group Address** – An IP multicast group address. If a group address is not specified, the RP is used for all multicast groups.
- ◆ **Group Mask** – Subnet mask that is used for the group address.

Web Interface

To configure a static rendezvous point:

1. Click Routing Protocol, PIM, SM.
2. Select RP Address from the Step list.
3. Specify the static RP to use for a multicast group, or a range of groups by using a subnet mask.
4. Click Apply.

Figure 420: Configuring a PIM Static Rendezvous Point

Routing Protocol > PIM > SM

Step: 3. RP Address Action: Add

RP Address: 192.168.1.1

Group Address: 224.9.0.0 (Optional)

Group Mask: 255.255.255.0 (Optional)

Apply Revert

To display static rendezvous points:

1. Click Routing Protocol, PIM, SM.
2. Select RP Address from the Step list.
3. Select Show from the Action list.

Figure 421: Showing PIM Static Rendezvous Points

Routing Protocol > PIM > SM

Step: 3. RP Address Action: Show

PIM-SM RP Address List Total: 1

<input type="checkbox"/>	RP Address	Group Address	Group Mask
<input type="checkbox"/>	192.168.1.1	224.9.0.0	255.255.255.0

Delete Revert

Configuring a PIM RP Candidate

Use the Routing Protocol > PIM > PIM-SM (RP Candidate) page to configure the switch to advertise itself as a Rendezvous Point (RP) candidate to the bootstrap router (BSR).

Command Usage

- ◆ When this router is configured as an RP candidate, it periodically sends PIMv2 messages to the BSR advertising itself as a candidate RP for the specified group addresses. The IP address of the designated VLAN is sent as the candidate's RP address. The BSR places information about all of the candidate RPs in subsequent bootstrap messages. The BSR uses the RP-election hash algorithm to select an active RP for each group range. The election process is performed by the BSR only for its own use. Each PIM-SM router that receives the list of RP candidates from the BSR also elects an active RP for each group range using the same election process.
- ◆ The election process for each group is based on the following criteria:
 - Find all RPs with the most specific group range.

- Select those with the highest priority (lowest priority value).
 - Compute hash value based on the group address, RP address, priority, and hash mask included in the bootstrap messages.
 - If there is a tie, use the candidate RP with the highest IP address.
- ◆ This distributed election process provides faster convergence and minimal disruption when an RP fails. It also serves to provide load balancing by distributing groups across multiple RPs. Moreover, when an RP fails, the responsible RPs are re-elected on each router, and the groups automatically distributed to the remaining RPs.
 - ◆ To improve failover recovery, it is advisable to select at least two core routers in diverse locations, each to serve as both a candidate BSR and candidate RP. It is also preferable to set up one of these routers as both the primary BSR and RP.

Parameters

These parameters are displayed:

- ◆ **VLAN** – Identifier of configured VLAN interface. (Range: 1-4094)
- ◆ **Interval** – The interval at which this device advertises itself as an RP candidate. (Range: 60-16383 seconds; Default: 60 seconds)
- ◆ **Priority** – Priority used by the candidate RP in the election process. The RP candidate with the largest priority is preferred. If the priority values are the same, the candidate with the larger IP address is elected to be the RP. Setting the priority to zero means that this router is not eligible to server as the RP. (Range: 0-255; Default: 0)
- ◆ **Group Address** – An IP multicast group address. If not defined, the default address is 224.0.0.0/4, or the entire IPv4 multicast group.
- ◆ **Group Mask** – Subnet mask that is used for the group address.

Web Interface

To advertise the switch as an RP candidate:

1. Click Multicast, Multicast Routing, SM.
2. Select RP Candidate from the Step list.
3. Specify a VLAN interface, the interval at which to advertise the router as an RP candidate, the priority to use in the election process, and the multicast group address and mask indicating the groups for which this router is bidding to become the RP.
4. Click Apply.

Figure 422: Configuring a PIM RP Candidate

Routing Protocol > PIM > SM

Step: 4. RP Candidate Action: Add

VLAN: 1

Interval (60-16383): 60 sec (Optional)

Priority (0-255): 100 (Optional)

Group Address: 224.0.0.0 (Optional)

Group Mask: 255.0.0.0 (Optional)

Note: If the group prefix is not defined, the default 224.0.0.0 240.0.0.0 is used.

Apply Revert

To display settings for an RP candidate:

1. Click Routing Protocol, PIM, SM.
2. Select RP Candidate from the Step list.
3. Select Show from the Action list.
4. Select an interface from the VLAN list.

Figure 423: Showing Settings for a PIM RP Candidate

Routing Protocol > PIM > SM

Step: 4. RP Candidate Action: Show

VLAN: 1

Interval: 60

Priority: 100

PIM-SM RP Candidate Group List Total: 1

Group Address	Group Mask
224.0.0.0	255.0.0.0

Delete

Displaying the PIM BSR Router Use the Routing Protocol > PIM > PIM-SM (Show Information – Show BSR Router) page to display Information about the bootstrap router (BSR).

Parameters

These parameters are displayed:

- ◆ **IP Address** – IP address of interface configured as the BSR.
- ◆ **Uptime** – The time this BSR has been up and running.
- ◆ **Priority** – Priority value used by this BSR candidate.

- ◆ **Hash Mask Length** – The number of significant bits used in the multicast group comparison mask by this BSR candidate.
- ◆ **Expire** – The time before the BSR is declared down.
- ◆ **Role** – Candidate or non-candidate BSR.
- ◆ **State**¹⁸ – Operation state of BSR includes:
 - No information – No information is stored for this device.
 - Accept Any – The router does not know of an active BSR, and will accept the first bootstrap message it sees as giving the new BSR's identity and the RP-set.
 - Accept Preferred – The router knows the identity of the current BSR, and is using the RP-set provided by that BSR. Only bootstrap messages from that BSR or from a C-BSR with higher weight than the current BSR will be accepted.
 - Candidate BSR – Bidding in election process.
 - Pending-BSR – The router is a candidate to be the BSR for the RP-set. Currently, no other router is the preferred BSR, but this router is not yet the elected BSR.
 - Elected BSR – Elected to serve as BSR.

Web Interface

To display information about the BSR:

1. Click Routing Protocol, PIM, SM.
2. Select Show Information from the Step list.
3. Select Show BSR Router from the Action list.

18. These parameters are based on RFC 5059.

Figure 424: Showing Information About the PIM BSR



The screenshot shows a web interface for configuring PIM. The breadcrumb is "Routing Protocol > PIM > SM". Below the breadcrumb, there are two dropdown menus: "Step:" with "5. Show Information" selected, and "Action:" with "Show BSR Router" selected. Below these are several parameters listed in a table-like format:

IP Address	192.168.0.2/32
Uptime	00:02:17
Priority	200
Hash Mask Length	20
Expire	00:00:07
Role	Candidate BSR
State	Elected BSR

Displaying PIM RP Mapping Use the Routing Protocol > PIM > PIM-SM (Show Information – Show RP Mapping) page to display active RPs and associated multicast routing entries.

Parameters

These parameters are displayed:

- ◆ **Groups** – A multicast group address.
- ◆ **RP Address** – IP address of the RP for the listed multicast group.
- ◆ **Information Source** – RP that advertised the mapping, how the RP was selected (Static or Bootstrap), and the priority used in the bidding process.
- ◆ **Uptime** – The time this RP has been up and running
- ◆ **Expire** – The time before this entry will be removed.

Web Interface

To display the RPs mapped to multicast groups:

1. Click Multicast, Multicast Routing, SM.
2. Select Show Information from the Step list.
3. Select Show RP Mapping from the Action list.

Figure 425: Showing PIM RP Mapping

The screenshot shows the 'Routing Protocol > PIM > SM' configuration page. At the top, there are two dropdown menus: 'Step: 5. Show Information' and 'Action: Show RP Mapping'. Below this is a table titled 'RP Mapping Information List' with a 'Total: 2' count. The table has five columns: Groups, RP Address, Information Source, Uptime, and Expire. There are two rows of data. Below the table is a 'Clear' button.

Groups	RP Address	Information Source	Uptime	Expire
172.16.0.0/16	10.6.6.6	10.6.6.6, via bootstrap, priority 0	22:36:49	00:02:04
192.168.0.0/24	10.9.9.9	10.9.9.9, via bootstrap, priority 0	22:36:20	00:03:27

Configuring PIMv6 for IPv6

This section describes how to configure PIM-DM and PIM-SM for IPv6.

Enabling PIMv6 Globally Use the Routing Protocol > PIM6 > General page to enable IPv6 PIM routing globally on the router.

Command Usage

- ◆ This feature enables PIM-DM and PIM-SM for IPv6 globally on the router. You also need to enable PIM-DM and PIM-SM for each interface that will support multicast routing (see [page 624](#)), and make any changes necessary to the multicast protocol parameters.
- ◆ To use PIMv6, multicast routing must be enabled on the switch (see [“Enabling Multicast Routing Globally” on page 602](#)).
- ◆ To use multicast routing, MLD proxy cannot be enabled on any interface of the device (see [“MLD Proxy Routing” in the CLI Reference Guide](#)).

Web Interface

To enable PIMv6 multicast routing:

1. Click Routing Protocol, PIM6, General.
2. Enable PIM6 Routing Protocol.
3. Click Apply.

Figure 426: Enabling PIMv6 Multicast Routing

The screenshot shows the 'Routing Protocol > PIM6 > General' configuration page. It features a checkbox labeled 'PIM6 Routing Protocol' which is checked and has the text 'Enabled' next to it. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

Configuring PIMv6 Interface Settings Use the Routing Protocol > PIM6 > Interface page configure the routing protocol's functional attributes for each interface.

Command Usage

- ◆ Most of the attributes on this page are common to both PIM6-DM and PIM6-SM. Select Dense or Sparse Mode to display the common attributes, as well as those applicable to the selected mode.
- ◆ An IPv6 address must first be assigned to the required routing interface before PIMv6 can be configured on this page.
- ◆ PIMv6 and MLD proxy cannot be used at the same time. When an interface is set to use PIMv6 Dense mode, MLD proxy cannot be enabled on any interface of the switch (see "MLD Proxy Routing" in the *CLI Reference Guide*). Also, when MLD proxy is enabled on an interface, PIMv6 cannot be enabled on any interface.

PIM6-DM

- ◆ PIM6-DM functions similar to DVMRP by periodically flooding the network with traffic from any active multicast server. It also uses MLD to determine the presence of multicast group members. The main difference, is that it uses the router's unicast routing table to determine if the interface through which a packet is received provides the shortest path back to the source.
- ◆ Dense-mode interfaces are subject to multicast flooding by default, and are only removed from the multicast routing table when the router determines that there are no group members or downstream routers, or when a prune message is received from a downstream router.

PIM6-SM

- ◆ A PIM6-SM interface is used to forward multicast traffic only if a join message is received from a downstream router or if group members are directly connected to the interface. When routers want to receive a multicast flow, they periodically send join messages to the RP, and are subsequently added to the shared path for the specified flow back up to the RP. If routers want to join the source path up through the SPT, they periodically send join messages toward the source. They also send prune messages toward the RP to prune the shared path once they have connected to the source through the SPT, or if there are no longer any group members connected to the interface.

Parameters

These parameters are displayed:

Common Attributes

- ◆ **VLAN** – Layer 3 VLAN interface. (Range: 1-4094)
- ◆ **Mode** – PIMv6 routing mode. (Options: Dense, None)
The routing mode must first be set to None, before changing between Dense and Sparse modes.
- ◆ **IPv6 Address** – IPv6 link-local address assigned to the selected VLAN.
- ◆ **Hello Holdtime** – Sets the interval to wait for hello messages from a neighboring PIM router before declaring it dead. Note that the hello holdtime should be greater than or equal to the value of Hello Interval, otherwise it will be automatically set to 3.5 x the Hello Interval. (Range: 1-65535 seconds; Default: 105 seconds, or 3.5 times the hello interval if set)
- ◆ **Hello Interval** – Sets the frequency at which PIM hello messages are transmitted out on all interfaces. (Range: 1-65535 seconds; Default: 30 seconds)
Hello messages are sent to neighboring PIM routers from which this device has received probes, and are used to verify whether or not these neighbors are still active members of the multicast tree. PIM-SM routers use these messages not only to inform neighboring routers of their presence, but also to determine which router for each LAN segment will serve as the Designated Router (DR).
When a router is booted or first configured to use PIM, it sends an initial hello message, and then sets its Hello timer to the configured value. If a router does not hear from a neighbor for the period specified by the Hello Holdtime, that neighbor is dropped. This hold time is included in each hello message received from a neighbor. Also note that hello messages also contain the DR priority of the router sending the message.
If the hello holdtime is already configured, and the hello interval is set to a value longer than the hello holdtime, this command will fail.
- ◆ **Join/Prune Holdtime** – Sets the hold time for the prune state. (Range: 1-65535 seconds; Default: 210 seconds)
 - PIM-DM: The multicast interface that first receives a multicast stream from a particular source forwards this traffic to all other PIM-DM interfaces on the router. If there are no requesting groups on that interface, the leaf node sends a prune message upstream and enters a prune state for this multicast stream. The prune state is maintained until the join/prune holdtime timer expires or a graft message is received for the forwarding entry.
 - PIM-SM: The multicast interface that first receives a multicast stream from a particular source forwards this traffic only to those interfaces on the router that have requests to join this group. When there are no longer any requesting groups on that interface, the leaf node sends a prune message upstream and enters a prune state for this multicast stream. The protocol

maintains both the current join state and the pending RPT prune state for this (source, group) pair until the join/prune interval timer expires.

- ◆ **LAN Prune Delay** – Causes this device to inform downstream routers of how long it will wait before pruning a flow after receiving a prune request. (Default: Disabled)

When other downstream routers on the same VLAN are notified that this upstream router has received a prune request, they must send a Join to override the prune before the prune delay expires if they want to continue receiving the flow. The message generated by this command effectively prompts any downstream neighbors with hosts receiving the flow to reply with a Join message. If no join messages are received after the prune delay expires, this router will prune the flow.

The sum of the Override Interval and Propagation Delay are used to calculate the LAN prune delay.

- ◆ **Override Interval** – The time required for a downstream router to respond to a LAN Prune Delay message by sending back a Join message if it wants to continue receiving the flow referenced in the message. (Range: 500-6000 milliseconds; Default: 2500 milliseconds)

The override interval and the propagation delay are used to calculate the LAN prune delay. If a downstream router has group members which want to continue receiving the flow referenced in a LAN prune delay message, then the override interval represents the time required for the downstream router to process the message and then respond by sending a Join message back to the upstream router to ensure that the flow is not terminated.

- ◆ **Propagation Delay** – The time required for a LAN prune delay message to reach downstream routers. (Range: 100-5000 milliseconds; Default: 500 milliseconds)

The override interval and propagation delay are used to calculate the LAN prune delay. If a downstream router has group members which want to continue receiving the flow referenced in a LAN prune delay message, then the propagation delay represents the time required for the LAN prune delay message to be propagated down from the upstream router to all downstream routers attached to the same VLAN interface.

- ◆ **Trigger Hello Delay** – The maximum time before transmitting a triggered PIM Hello message after the router is rebooted or PIM is enabled on an interface. (Range: 0-5 seconds; Default: 5 seconds)

When a router first starts or PIM is enabled on an interface, the hello delay is set to random value between 0 and the trigger hello delay. This prevents synchronization of Hello messages on multi-access links if multiple routers are powered on simultaneously.

Also, if a Hello message is received from a new neighbor, the receiving router will send its own Hello message after a random delay between 0 and the trigger hello delay.

Dense-Mode Attributes

- ◆ **Graft Retry Interval** – The time to wait for a Graft acknowledgement before resending a Graft message. (Range: 1-10 seconds; Default: 3 seconds)

A graft message is sent by a router to cancel a prune state. When a router receives a graft message, it must respond with an graft acknowledgement message. If this acknowledgement message is lost, the router that sent the graft message will resend it a number of times (as defined by Max. Graft Retries).

- ◆ **Max. Graft Retries** – The maximum number of times to resend a Graft message if it has not been acknowledged. (Range: 1-10; Default: 3)

- ◆ **State Refresh Origination Interval** – The interval between sending PIM-DM state refresh control messages. (Range: 1-100 seconds; Default: 60 seconds)

The pruned state times out approximately every three minutes and the entire PIM-DM network is re-flooded with multicast packets and prune messages. The state refresh feature keeps the pruned state from timing out by periodically forwarding a control message down the distribution tree, refreshing the prune state on the outgoing interfaces of each router in the tree. This also enables PIM routers to recognize topology changes (sources joining or leaving a multicast group) before the default three-minute state timeout expires.

This command is only effectively for interfaces of first hop, PIM-DM routers that are directly connected to the sources of multicast groups.

Sparse-Mode Attributes

- ◆ **DR Priority** – Sets the priority advertised by a router when bidding to become the Designated Router (DR). (Range: 0-4294967294; Default: 1)

More than one PIM-SM router may be connected to an Ethernet or other shared-media LAN. If multicast hosts are directly connected to the LAN, then only one of these routers is elected as the DR, and acts on behalf of these hosts, sending periodic Join/Prune messages toward a group-specific RP for each group. A single DR is elected per interface (LAN or otherwise) using a simple election process.

The router with the highest priority configured on an interface is elected as the DR. If more than one router attached to this interface uses the same priority, then the router with the highest IP address is elected to serve as the DR.

If a router does not advertise a priority in its hello messages, it is assumed to have the highest priority and is elected as the DR. If more than one router is not advertising its priority, then the router with the highest IP address is elected to serve as the DR.

- ◆ **Join/Prune Interval** – Sets the interval at which join/prune messages are sent. (Range: 1-65535 seconds; Default: 60 seconds)

By default, the switch sends join/prune messages every 60 seconds to inform other PIM-SM routers about clients who want to join or leave a multicast group.

Use the same join/prune message interval on all PIM-SM routers in the same PIM-SM domain, otherwise the routing protocol's performance will be adversely affected.

The multicast interface that first receives a multicast stream from a particular source forwards this traffic only to those interfaces on the router that have requests to join this group. When there are no longer any requesting groups on that interface, the leaf node sends a prune message upstream and enters a prune state for this multicast stream. The protocol maintains both the current join state and the pending RPT prune state for this (source, group) pair until the join/prune interval timer expires.

Web Interface

To configure PIMv6 interface settings:

1. Click Routing Protocol, PIM6, Interface.
2. Modify any of the protocol parameters as required.
3. Click Apply.

Figure 427: Configuring PIMv6 Interface Settings (Dense Mode)

The screenshot shows a configuration page titled "Routing Protocol > PIM6 > Interface". The page contains the following settings:

VLAN	1
Mode	Dense
IPv6 Address	FE80::200:E8FF:FE90:0
Hello Holdtime (1-65535)	105 sec
Hello Interval (1-65535)	30 sec
Join/Prune Holdtime (1-65535)	210 sec
LAN Prune Delay	<input type="checkbox"/> Enabled
Override Interval (500-6000)	2500 msec
Propagation Delay (100-5000)	500 msec
Trigger Hello Delay (0-5)	5 sec
Graft Retry Interval (1-10)	3 sec
Max. Graft Retries (1-10)	3
State Refresh Origination Interval (1-100)	60 sec

At the bottom right of the form, there are two buttons: "Apply" and "Revert".

Figure 428: Configuring PIMv6 Interface Settings (Sparse Mode)

Routing Protocol > PIM6 > Interface	
VLAN	1
Mode	Sparse
IPv6 Address	FE80::200:CFF:FE00:FD
Hello Holdtime (1-65535)	105 sec
Hello Interval (1-65535)	30 sec
Join/Prune Holdtime (1-65535)	210 sec
LAN Prune Delay	<input type="checkbox"/> Enabled
Override Interval (500-6000)	2500 msec
Propagation Delay (100-5000)	500 msec
Trigger Hello Delay (0-5)	5 sec
DR Priority (0-4294967294)	1
Join/Prune Interval (1-65535)	60 sec

Displaying PIM6 Neighbor Information Use the Routing Protocol > PIM6 > Neighbor page to display all neighboring PIMv6 routers.

Parameters

These parameters are displayed:

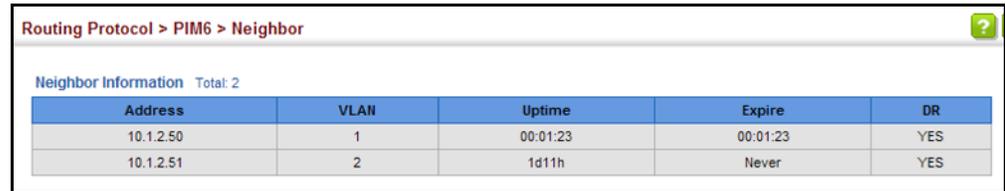
- ◆ **Address** – IP address of the next-hop router.
- ◆ **VLAN** – VLAN that is attached to this neighbor.
- ◆ **Uptime** – The duration this entry has been active.
- ◆ **Expire** – The time before this entry will be removed.
- ◆ **DR** – The designated PIM6-SM router. If multicast hosts are directly connected to the LAN, then only one of these routers is elected as the DR, and acts on behalf of these hosts, sending periodic Join/Prune messages toward a group-specific RP for each group.

Web Interface

To display neighboring PIMv6 routers:

1. Click Routing Protocol, PIM6, Neighbor.

Figure 429: Showing PIMv6 Neighbors



Address	VLAN	Uptime	Expire	DR
10.1.2.50	1	00:01:23	00:01:23	YES
10.1.2.51	2	1d11h	Never	YES

Configuring Global PIM6-SM Settings

Use the Routing Protocol > PIM6 > PIM6-SM (Configure Global) page to configure the rate at which register messages are sent, the source of register messages, and switch over to the Shortest Path Tree (SPT).

Parameters

These parameters are displayed:

- ◆ **Register Rate Limit** – Configures the rate at which register messages are sent by the Designated Router (DR) for each (source, group) entry. (Range: 1-65535 packets per second; Default: disabled)

This parameter can be used to relieve the load on the designated router (DR) and rendezvous point (RP). However, because register messages exceeding the limit are dropped, some receivers may experience data packet loss within the first few seconds in which register messages are sent from bursty sources.

- ◆ **Register Source** – Configures the IP source address of a register message to an address other than the outgoing interface address of the DR that leads back toward the RP. (Range: VLAN 1-4094; Default: The IP address of the DR's outgoing interface that leads back to the RP)

When the source address of a register message is filtered by intermediate network devices, or is not a uniquely routed address to which the RP can send packets, the replies sent from the RP to the source address will fail to reach the DR, resulting in PIM6-SM protocol failures. This type of problem can be overcome by manually configuring the source address of register messages to an interface that leads back to the RP.

- ◆ **SPT Threshold** – Prevents the last-hop PIM-SM router from switching to Shortest Path Source Tree (SPT) mode. (Options: Infinity, Reset; Default: Reset)

The default path for packets from a multicast source to a receiver is through the RP. However, the path through the RP is not always the shortest path. Therefore, the router uses the RP to forward only the first packet from a new multicast group to its receivers. Afterwards, it calculates the shortest path tree (SPT) directly between the receiver and source, and then uses the SPT to send all subsequent packets from the source to the receiver instead of using the shared

tree. Note that when the SPT threshold is not set by this command, the PIM leaf router will join the shortest path tree immediately after receiving the first packet from a new source.

Enable the SPT threshold to force the router to use the shared tree for all multicast groups, or just for the specified multicast groups.

- ◆ **Group Address** – An IPv6 multicast group address. If a group address is not specified, the shared tree is used for all multicast groups.
- ◆ **Group Prefix Length** – An IPv6 network prefix length for a multicast group. (Range: 8-128)

Web Interface

To configure global settings for PIM6-SM:

1. Click Routing Protocol, PIM6, SM.
2. Select Configure Global from the Step list.
3. Set the register rate limit and source of register messages if required. Also specify any multicast groups which must be routed across the shared tree, instead of switching over to the SPT.
4. Click Apply.

Figure 430: Configuring Global Settings for PIM6-SM

The screenshot shows a web configuration page for PIM6-SM. The breadcrumb is 'Routing Protocol > PIM6 > SM'. The 'Step' dropdown is '1. Configure Global'. The configuration fields are:

- Register Rate Limit (1-65535): Enabled, 1500 packets/sec
- Register Source: Enabled, VLAN 1
- SPT Threshold: Infinity
- Group Address: FF00:: (Optional)
- Group Prefix Length (8-128): 8 (Optional)

Buttons: Apply, Revert

Configuring a PIM6 BSR Candidate Use the Routing Protocol > PIM6 > PIM6-SM (BSR Candidate) page to configure the switch as a Bootstrap Router (BSR) candidate.

Command Usage

- ◆ When this router is configured as a BSR candidate, it starts sending bootstrap messages to all of its PIM6-SM neighbors. The primary IP address of the designated VLAN is sent as the candidate's BSR address. Each neighbor receiving the bootstrap message compares the BSR address with the address from previous messages. If the current address is the same or a higher address,

it accepts the bootstrap message and forwards it. Otherwise, it drops the message.

- ◆ This router will continue to be the BSR until it receives a bootstrap message from another candidate with a higher priority (or a higher IP address if the priorities are the same).
- ◆ To improve failover recovery, it is advisable to select at least two core routers in diverse locations, each to serve as both a candidate BSR and candidate RP. It is also preferable to set up one of these routers as both the primary BSR and RP.

Parameters

These parameters are displayed:

- ◆ **BSR Candidate Status** – Configures the switch as a Bootstrap Router (BSR) candidate. (Default: Disabled)
- ◆ **VLAN ID** – Identifier of configured VLAN interface. (Range: 1-4093)
- ◆ **Hash Mask Length** – Hash mask length (in bits) used for RP selection (see [“Configuring a PIM6 Static Rendezvous Point” on page 633](#) and [“Configuring a PIM6 RP Candidate” on page 635](#)). The portion of the hash specified by the mask length is ANDed with the group address. Therefore, when the hash function is executed on any BSR, all groups with the same seed hash will be mapped to the same RP. If the mask length is less than 32, then only the first portion of the hash is used, and a single RP will be defined for multiple groups. (Range: 0-32; Default: 10)
- ◆ **Priority** – Priority used by the candidate bootstrap router in the election process. The BSR candidate with the largest priority is preferred. If the priority values are the same, the candidate with the larger IP address is elected to be the BSR. Setting the priority to zero means that this router is not eligible to server as the BSR. At least one router in the PIM6-SM domain must be set to a value greater than zero. (Range: 0-255; Default: 0)

Web Interface

To configure the switch as a BSR candidate:

1. Click Routing Protocol, PIM6, SM.
2. Select BSR Candidate from the Step list.
3. Specify the VLAN interface for which this router is bidding to become the BSR, the hash mask length that will subsequently be used for RP selection if this router is selected as the BSR, and the priority for BSR selection.
4. Click Apply.

Figure 431: Configuring a PIM6-SM BSR Candidate

Configuring a PIM6 Static Rendezvous Point

Use the Routing Protocol > PIM6 > PIM6-SM (RP Address) page to configure a static address as the Rendezvous Point (RP) for a particular multicast group.

Command Usage

- ◆ The router will act as an RP for all multicast groups in the local PIM6-SM domain if no groups are specified. A static RP can either be configured for the whole multicast group range FF00::/8, or for specific group ranges.
- ◆ If an IP address is specified that was previously used for an RP, then the older entry is replaced.
- ◆ Multiple RPs can be defined for different groups or group ranges. If a group is matched by more than one entry, the router will use the RP associated with the longer group prefix length. If the prefix lengths are the same, then the static RP with the highest IP address is chosen.
- ◆ Static definitions for RP addresses may be used together with RP addresses dynamically learned through the bootstrap router (BSR). If an RP address learned by the BSR and one statically configured using this command are both available for a group range, the RP address learned by the BSR is chosen over the one statically configured.
- ◆ All routers within the same PIM6-SM domain must be configured with the same RP(s). Selecting an RP through the dynamic election process is therefore preferable for most situations. Using the dynamic RP election process also allows a backup RP to automatically take over if the active RP router becomes unavailable.

Parameters

These parameters are displayed:

- ◆ **RP Address** – Static IP address of the router that will be an RP for the specified multicast group(s).
- ◆ **Group Address** – An IP multicast group address. If a group address is not specified, the RP is used for all multicast groups.

- ◆ **Group Prefix Length** – An IPv6 network prefix length for a multicast group. (Range: 8-128)

Web Interface

To configure a static rendezvous point:

1. Click Routing Protocol, PIM6, SM.
2. Select RP Address from the Step list.
3. Specify the static RP to use for a multicast group, or a range of groups by using a subnet mask.
4. Click Apply.

Figure 432: Configuring a PIM6 Static Rendezvous Point

Routing Protocol > PIM6 > SM

Step: 3. RP Address Action: Add

RP Address: 80::200:CFF:FE00:FD

Group Address: (Optional)

Group Prefix Length (8-128): (Optional)

Apply Revert

To display static rendezvous points:

1. Click Routing Protocol, PIM6, PIM6-SM.
2. Select RP Address from the Step list.
3. Select Show from the Action list.

Figure 433: Showing PIM6 Static Rendezvous Points

Routing Protocol > PIM6 > SM

Step: 3. RP Address Action: Show

PIM6-SM RP Address List Total: 1

	RP Address	Group Address
<input type="checkbox"/>	FE80::200:E8FF:FE93:82A0/128	FF00::8

Delete Revert

Configuring a PIM6 RP Candidate Use the Routing Protocol > PIM6 > PIM6-SM (RP Candidate) page to configure the switch to advertise itself as a Rendezvous Point (RP) candidate to the bootstrap router (BSR).

Command Usage

- ◆ When this router is configured as an RP candidate, it periodically sends PIMv2 messages to the BSR advertising itself as a candidate RP for the specified group addresses. The IP address of the designated VLAN is sent as the candidate's RP address. The BSR places information about all of the candidate RPs in subsequent bootstrap messages. The BSR uses the RP-election hash algorithm to select an active RP for each group range. The election process is performed by the BSR only for its own use. Each PIM6-SM router that receives the list of RP candidates from the BSR also elects an active RP for each group range using the same election process.
- ◆ The election process for each group is based on the following criteria:
 - Find all RPs with the most specific group range.
 - Select those with the highest priority (lowest priority value).
 - Compute hash value based on the group address, RP address, priority, and hash mask included in the bootstrap messages.
 - If there is a tie, use the candidate RP with the highest IP address.
- ◆ This distributed election process provides faster convergence and minimal disruption when an RP fails. It also serves to provide load balancing by distributing groups across multiple RPs. Moreover, when an RP fails, the responsible RPs are re-elected on each router, and the groups automatically distributed to the remaining RPs.
- ◆ To improve failover recovery, it is advisable to select at least two core routers in diverse locations, each to serve as both a candidate BSR and candidate RP. It is also preferable to set up one of these routers as both the primary BSR and RP.

Parameters

These parameters are displayed:

- ◆ **VLAN** – Identifier of configured VLAN interface. (Range: 1-4093)
- ◆ **Interval** – The interval at which this device advertises itself as an RP candidate. (Range: 60-16383 seconds; Default: 60 seconds)
- ◆ **Priority** – Priority used by the candidate RP in the election process. The RP candidate with the largest priority is preferred. If the priority values are the same, the candidate with the larger IP address is elected to be the RP. Setting the priority to zero means that this router is not eligible to server as the RP. (Range: 0-255; Default: 0)

- ◆ **Group Address** – An IP multicast group address. If not defined, the RP is advertised for all multicast groups.
- ◆ **Group Prefix Length** – Subnet mask that is used for the group address. (Range: 8-128)

Web Interface

To advertise the switch as an RP candidate:

1. Click Routing Protocol, PIM6, SM.
2. Select RP Candidate from the Step list.
3. Specify a VLAN interface, the interval at which to advertise the router as an RP candidate, the priority to use in the election process, and the multicast group address and mask indicating the groups for which this router is bidding to become the RP.
4. Click Apply.

Figure 434: Configuring a PIM6 RP Candidate

Routing Protocol > PIM6 > SM	
Step:	4. RP Candidate
Action:	Add
VLAN ID	1
Interval (60-16383)	sec (Optional)
Priority (0-255)	(Optional)
Group Address	FFAA::0101 (Optional)
Group Prefix Length (8-128)	8 (Optional)
<input type="button" value="Apply"/> <input type="button" value="Revert"/>	

To display settings for an RP candidate:

1. Click Routing Protocol, PIM6, SM.
2. Select RP Candidate from the Step list.
3. Select Show from the Action list.
4. Select an interface from the VLAN list.

Figure 435: Showing Settings for a PIM6 RP Candidate

Routing Protocol > PIM6 > SM				
Step: 4. RP Candidate		Action: Show		
PIM6-SM RP Candidate List Total: 1				
<input type="checkbox"/>	VLAN	Interval	Priority	Group Address / Prefix Length
<input type="checkbox"/>	1	60	0	FF00::8

Delete Revert

Displaying the PIM6 BSR Router Use the Routing Protocol > PIM6 > PIM6-SM (Show Information – Show BSR Router) page to display Information about the bootstrap router (BSR).

Parameters

These parameters are displayed:

- ◆ **IP Address** – IP address of interface configured as the BSR.
- ◆ **Uptime** – The time this BSR has been up and running.
- ◆ **Priority** – Priority value used by this BSR candidate.
- ◆ **Hash Mask Length** – The number of significant bits used in the multicast group comparison mask by this BSR candidate.
- ◆ **Expire** – The time before the BSR is declared down.
- ◆ **Role** – Candidate or non-candidate BSR.
- ◆ **State**¹⁹ – Operation state of BSR includes:
 - No information – No information is stored for this device.
 - Accept Any – The router does not know of an active BSR, and will accept the first bootstrap message it sees as giving the new BSR's identity and the RP-set.
 - Accept Preferred – The router knows the identity of the current BSR, and is using the RP-set provided by that BSR. Only bootstrap messages from that BSR or from a C-BSR with higher weight than the current BSR will be accepted.
 - Candidate BSR – Bidding in election process.
 - Pending-BSR – The router is a candidate to be the BSR for the RP-set. Currently, no other router is the preferred BSR, but this router is not yet the elected BSR.

19. These parameters are based on RFC 5059.

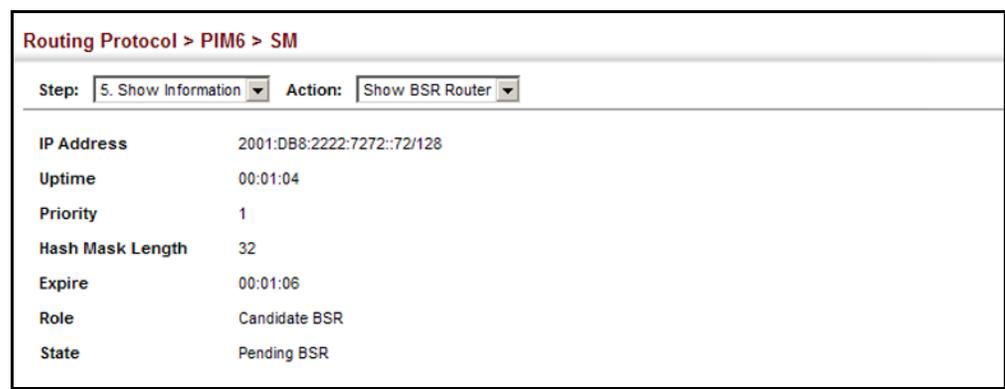
- Elected BSR – Elected to serve as BSR.

Web Interface

To display information about the BSR:

1. Click Routing Protocol, PIM6, SM.
2. Select Show Information from the Step list.
3. Select Show BSR Router from the Action list.

Figure 436: Showing Information About the PIM6 BSR



The screenshot shows a web interface for configuring PIM6. The breadcrumb path is "Routing Protocol > PIM6 > SM". Below the breadcrumb, there are two dropdown menus: "Step: 5. Show Information" and "Action: Show BSR Router". The main content area displays the following information:

IP Address	2001:DB8:2222:7272::72/128
Uptime	00:01:04
Priority	1
Hash Mask Length	32
Expire	00:01:06
Role	Candidate BSR
State	Pending BSR

Displaying RP Mapping Use the Routing Protocol > PIM6 > PIM6-SM (Show Information – Show RP Mapping) page to display active RPs and associated multicast routing entries.

Parameters

These parameters are displayed:

- ◆ **Groups** – A multicast group address.
- ◆ **RP Address** – IP address of the RP for the listed multicast group.
- ◆ **Information Source** – RP that advertised the mapping, how the RP was selected (Static or Bootstrap), and the priority used in the bidding process.
- ◆ **Uptime** – The time this RP has been up and running
- ◆ **Expire** – The time before this entry will be removed.

Web Interface

To display the RPs mapped to multicast groups:

1. Click Routing Protocol, PIM6, SM.
2. Select Show Information from the Step list.
3. Select Show RP Mapping from the Action list.

Figure 437: Showing PIM6 RP Mapping

Routing Protocol > PIM6 > SM

Step: 5. Show Information Action: Show RP Mapping

RP Mapping Information List Total: 1

Groups	RP Address	Information Source	Uptime	Expire
FF00::8	FE80::200:CFF:FE00:FD/128	static	00:24:22	Never

Clear

Section III

Appendices

This section provides additional information and includes these items:

- ◆ [“Software Specifications” on page 643](#)
- ◆ [“Troubleshooting” on page 649](#)
- ◆ [“License Information” on page 651](#)



Software Specifications

Software Features

Management Authentication Local, RADIUS, TACACS+, Port Authentication (802.1X), HTTPS, SSH, Port Security, IP Filter

Client Access Control Access Control Lists (2048 rules), Port Authentication (802.1X), MAC Authentication, Port Security, DHCP Snooping, IP Source Guard

Port Configuration 1000BASE-SX/LX - 1000 Mbps full duplex (SFP)
10GBASE-CR/SR/LR/LRM - 10 Gbps full duplex (SFP+)
40GBASET-CR4 - 40 Gbps full duplex (QSFP+)

Flow Control Full Duplex: IEEE 802.3-2005
Half Duplex: Back pressure

Storm Control Broadcast, multicast, or unknown unicast traffic throttled above a critical threshold

Port Mirroring 2 sessions, one or more source ports to one destination port

Rate Limits Input/Output Limits
Range configured per port

Port Trunking Static trunks (Cisco EtherChannel compliant)
Dynamic trunks (Link Aggregation Control Protocol)

Spanning Tree Algorithm Spanning Tree Protocol (STP, IEEE 802.1D-2004)
Rapid Spanning Tree Protocol (RSTP, IEEE 802.1D-2004)
Multiple Spanning Tree Protocol (MSTP, IEEE 802.1D-2004)

VLAN Support Up to 4094 groups; port-based, tagged (802.1Q), QinQ tunnel,

Class of Service Supports eight levels of priority
Strict, Weighted Round Robin (WRR), or combination of strict and weighted queueing
Layer 3/4 priority mapping: IP Port, IP Precedence, IP DSCP

Quality of Service DiffServ supports class maps, policy maps, and service policies

Multicast Filtering IGMP Snooping (Layer 2 IPv4)
MLD Snooping (Layer 2 IPv6)
IGMP (Layer 3)
Multicast VLAN Registration (IPv4/IPv6)

IP Routing ARP, Proxy ARP
Static routes
CIDR (Classless Inter-Domain Routing)
RIP, RIPv2, OSPFv2, OSPFv3 unicast routing
PIM-SM, PIM-DM, PIMv6 multicast routing
VRRP (Virtual Router Redundancy Protocol)

Additional Features BOOTP Client,
Connectivity Fault Management
DHCP Client, Relay, Option 82, Server
DNS Client, Proxy
LLDP (Link Layer Discover Protocol)
RMON (Remote Monitoring, groups 1,2,3,9)
SMTP Email Alerts
SNMP (Simple Network Management Protocol)
SNTP (Simple Network Time Protocol)

Management Features

In-Band Management Telnet, web-based HTTP or HTTPS, SNMP manager, or Secure Shell

Out-of-Band Management RS-232 DB-9 console port

Software Loading HTTP, FTP or TFTP in-band, or XModem out-of-band

SNMP Management access via MIB database
Trap management to specified hosts

RMON Groups 1, 2, 3, 9 (Statistics, History, Alarm, Event)

Standards

BGPv4 (RFC 4271)
 IEEE 802.1AB Link Layer Discovery Protocol
 IEEE 802.1D-2004 Spanning Tree Algorithm and traffic priorities
 Spanning Tree Protocol
 Rapid Spanning Tree Protocol
 Multiple Spanning Tree Protocol
 IEEE 802.1p Priority tags
 IEEE 802.1Q VLAN
 IEEE 802.1v Protocol-based VLANs
 IEEE 802.1X Port Authentication
 IEEE 802.3-2005
 Ethernet, Fast Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet
 Link Aggregation Control Protocol (LACP)
 Full-duplex flow control (ISO/IEC 8802-3)
 IEEE 802.3ba-2010
 IEEE 802.3ac VLAN tagging
 IEEE 802.1ag Connectivity Fault Management (Amendment 5, D7.1)
 ARP (RFC 826)
 DHCP Client (RFC 2131)
 DHCP Relay (RFC 951, 2132, 3046)
 DHCP Server (RFC 2131, 2132)
 HTTPS
 ICMP (RFC 792)
 IGMP (RFC 1112)
 IGMPv2 (RFC 2236)
 IGMPv3 (RFC 3376) - partial support
 IGMP Proxy (RFC 4541)
 IPv4 IGMP (RFC 3228)
 MLD Snooping (RFC 4541)
 NTP (RFC 1305)
 OSPF (RFC 2328, 2178, 1587)
 OSPFv3 (RFC 2740)
 PIM-SM (RFC 4601)

PIM-DM (RFC 3973)
RADIUS+ (RFC 2618)
RIPv1 (RFC 1058)
RIPv2 (RFC 2453)
RIPv2, extension (RFC 1724)
RMON (RFC 2819 groups 1,2,3,9)
SNMP (RFC 1157)
SNMPv2c (RFC 1901, 2571)
SNMPv3 (RFC DRAFT 2273, 2576, 3410, 3411, 3413, 3414, 3415)
SNTP (RFC 2030)
SSH (Version 2.0)
TELNET (RFC 854, 855, 856)
TFTP (RFC 1350)
VRRP (RFC 3768)

Management Information Bases

Bridge MIB (RFC 1493)
Differentiated Services MIB (RFC 3289)
DNS Resolver MIB (RFC 1612)
Entity MIB (RFC 2737)
Ether-like MIB (RFC 2665)
Extended Bridge MIB (RFC 2674)
Extensible SNMP Agents MIB (RFC 2742)
Forwarding Table MIB (RFC 2096)
IGMP MIB (RFC 2933)
Interface Group MIB (RFC 2233)
Interfaces Evolution MIB (RFC 2863)
IP MIB (RFC 2011)
IP Forwarding Table MIB (RFC 2096)
IP Multicasting related MIBs
IPV6-MIB (RFC 2065)
IPV6-ICMP-MIB (RFC 2066)
IPV6-TCP-MIB (RFC 2052)
IPV6-UDP-MIB (RFC2054)
Link Aggregation MIB (IEEE 802.3ad)
MAU MIB (RFC 3636)
MIB II (RFC 1213)
OSPF MIB (RFC 1850)
OSPFv3 MIB (draft-ietf-ospf-ospfv3-mib-15.txt)
P-Bridge MIB (RFC 2674P)
Port Access Entity MIB (IEEE 802.1X)
Port Access Entity Equipment MIB

Power Ethernet MIB (RFC 3621)
Private MIB
Q-Bridge MIB (RFC 2674Q)
QinQ Tunneling (IEEE 802.1ad Provider Bridges)
Quality of Service MIB
RADIUS Accounting Server MIB (RFC 2621)
RADIUS Authentication Client MIB (RFC 2619)
RIP1 MIB (RFC 1058)
RIP2 MIB (RFC 2453)
RIP2 Extension (RFC1724)
RMON MIB (RFC 2819)
RMON II Probe Configuration Group (RFC 2021, partial implementation)
SNMP Community MIB (RFC 3584)
SNMP Framework MIB (RFC 3411)
SNMP-MPD MIB (RFC 3412)
SNMP Target MIB, SNMP Notification MIB (RFC 3413)
SNMP User-Based SM MIB (RFC 3414)
SNMP View Based ACM MIB (RFC 3415)
SNMPv2 IP MIB (RFC 2011)
TACACS+ Authentication Client MIB
TCP MIB (RFC 2012)
Trap (RFC 1215)
UDP MIB (RFC 2013)
VRRP MIB (RFC 2787)



Troubleshooting

Problems Accessing the Management Interface

Table 45: Troubleshooting Chart

Symptom	Action
Cannot connect using a web browser	<ul style="list-style-type: none">◆ Be sure the switch is powered on.◆ Check network cabling between the management station and the switch. Make sure the ends are properly connected and there is no damage to the cable. Test the cable if necessary.◆ Check that you have a valid network connection to the switch and that the port you are using has not been disabled.◆ Be sure you have configured the VLAN interface through which the management station is connected with a valid IP address, subnet mask and default gateway.◆ Be sure the management station has an IP address in the same subnet as the switch's IP interface to which it is connected.◆ If you are trying to connect to the switch via the IP address for a tagged VLAN group, your management station, and the ports connecting intermediate switches in the network, must be configured with the appropriate tag.
Cannot access the on-board configuration program via a serial port connection	<ul style="list-style-type: none">◆ Refer to the <i>CLI Reference Guide</i> for information on troubleshooting a connection to the serial port
Forgot or lost the password	<ul style="list-style-type: none">◆ Contact your local distributor.

Using System Logs

If a fault does occur, refer to the Installation Guide to ensure that the problem you encountered is actually caused by the switch. If the problem appears to be caused by the switch, follow these steps:

1. Enable logging.
2. Set the error messages reported to include all categories.
3. Enable SNMP.
4. Enable SNMP traps.
5. Designate the SNMP host that is to receive the error messages.
6. Repeat the sequence of commands or other actions that lead up to the error.
7. Make a list of the commands or circumstances that led to the fault. Also make a list of any error messages displayed.
8. Set up your terminal emulation software so that it can capture all console output to a file. Then enter the “show tech-support” command to record all system settings in this file.
9. Contact your distributor’s service engineer, and send a detailed description of the problem, along with the file used to record your system settings.

For example:

```
Console(config)#logging on
Console(config)#logging history flash 7
Console(config)#snmp-server host 192.168.1.23
:
```



License Information

This product includes copyrighted third-party software subject to the terms of the GNU General Public License (GPL), GNU Lesser General Public License (LGPL), or other related free software licenses. The GPL code used in this product is distributed WITHOUT ANY WARRANTY and is subject to the copyrights of one or more authors. For details, refer to the section "The GNU General Public License" below, or refer to the applicable license as included in the source-code archive.

The GNU General Public License

GNU GENERAL PUBLIC LICENSE
Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

1. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

2. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

3. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
 - b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
 - c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

4. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 - a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

- 5. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
- 6. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
- 7. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
- 8. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

9. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
10. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.
11. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

1. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
2. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Glossary

ACL Access Control List. ACLs can limit network traffic and restrict access to certain users or devices by checking each packet for certain IP or MAC (i.e., Layer 2) information.

ARP Address Resolution Protocol converts between IP addresses and MAC (hardware) addresses. ARP is used to locate the MAC address corresponding to a given IP address. This allows the switch to use IP addresses for routing decisions and the corresponding MAC addresses to forward packets from one hop to the next.

CoS Class of Service is supported by prioritizing packets based on the required level of service, and then placing them in the appropriate output queue. Data is transmitted from the queues using weighted round-robin service to enforce priority service and prevent blockage of lower-level queues. Priority may be set according to the port default, the packet's priority bit (in the VLAN tag), TCP/UDP port number, IP Precedence bit, or DSCP priority bit.

DHCP Dynamic Host Control Protocol. Provides a framework for passing configuration information to hosts on a TCP/IP network. DHCP is based on the Bootstrap Protocol (BOOTP), adding the capability of automatic allocation of reusable network addresses and additional configuration options.

DHCP Option 82 A relay option for sending information about the requesting client (or an intermediate relay agent) in the DHCP request packets forwarded by the switch and in reply packets sent back from the DHCP server. This information can be used by DHCP servers to assign fixed IP addresses, or set other services or policies for clients.

DHCP Snooping A technique used to enhance network security by snooping on DHCP server messages to track the physical location of hosts, ensure that hosts only use the IP addresses assigned to them, and ensure that only authorized DHCP servers are accessible.

DiffServ Differentiated Services provides quality of service on large networks by employing a well-defined set of building blocks from which a variety of aggregate forwarding behaviors may be built. Each packet carries information (DS byte) used by each hop to give it a particular forwarding treatment, or per-hop behavior, at each network node. DiffServ allocates different levels of service to users on the network with mechanisms such as traffic meters, shapers/droppers, packet markers at the boundaries of the network.

- DNS** Domain Name Service. A system used for translating host names for network nodes into IP addresses.
- DSCP** Differentiated Services Code Point Service. DSCP uses a six-bit tag to provide for up to 64 different forwarding behaviors. Based on network policies, different kinds of traffic can be marked for different kinds of forwarding. The DSCP bits are mapped to the Class of Service categories, and then into the output queues.
- EAPOL** Extensible Authentication Protocol over LAN. EAPOL is a client authentication protocol used by this switch to verify the network access rights for any device that is plugged into the switch. A user name and password is requested by the switch, and then passed to an authentication server (e.g., RADIUS) for verification. EAPOL is implemented as part of the IEEE 802.1X Port Authentication standard.
- EUI** Extended Universal Identifier is an address format used by IPv6 to identify the host portion of the network address. The interface identifier in EUI compatible addresses is based on the link-layer (MAC) address of an interface. Interface identifiers used in global unicast and other IPv6 address types are 64 bits long and may be constructed in the EUI-64 format. The modified EUI-64 format interface ID is derived from a 48-bit link-layer address by inserting the hexadecimal number FFFE between the upper three bytes (OUI field) and the lower 3 bytes (serial number) of the link layer address. To ensure that the chosen address is from a unique Ethernet MAC address, the 7th bit in the high-order byte is set to 1 (equivalent to the IEEE Global/Local bit) to indicate the uniqueness of the 48-bit address.
- GARP** Generic Attribute Registration Protocol. GARP is a protocol that can be used by endstations and switches to register and propagate multicast group membership information in a switched environment so that multicast data frames are propagated only to those parts of a switched LAN containing registered endstations. Formerly called Group Address Registration Protocol.
- GMRP** Generic Multicast Registration Protocol. GMRP allows network devices to register end stations with multicast groups. GMRP requires that any participating network devices or end stations comply with the IEEE 802.1p standard.
- GVRP** GARP VLAN Registration Protocol. Defines a way for switches to exchange VLAN information in order to register necessary VLAN members on ports along the Spanning Tree so that VLANs defined in each switch can work automatically over a Spanning Tree network.
- ICMP** Internet Control Message Protocol is a network layer protocol that reports errors in processing IP packets. ICMP is also used by routers to feed back information about better routing choices.

- IEEE 802.1D** Specifies a general method for the operation of MAC bridges, including the Spanning Tree Protocol.
- IEEE 802.1Q** VLAN Tagging—Defines Ethernet frame tags which carry VLAN information. It allows switches to assign endstations to different virtual LANs, and defines a standard way for VLANs to communicate across switched networks.
- IEEE 802.1p** An IEEE standard for providing quality of service (QoS) in Ethernet networks. The standard uses packet tags that define up to eight traffic classes and allows switches to transmit packets based on the tagged priority value.
- IEEE 802.1s** An IEEE standard for the Multiple Spanning Tree Protocol (MSTP) which provides independent spanning trees for VLAN groups.
- IEEE 802.1w** An IEEE standard for the Rapid Spanning Tree Protocol (RSTP) which reduces the convergence time for network topology changes to about 10% of that required by the older IEEE 802.1D STP standard. (Now incorporated in IEEE 802.1D-2004)
- IEEE 802.1X** Port Authentication controls access to the switch ports by requiring users to first enter a user ID and password for authentication.
- IEEE 802.3ac** Defines frame extensions for VLAN tagging.
- IEEE 802.3x** Defines Ethernet frame start/stop requests and timers used for flow control on full-duplex links. (Now incorporated in IEEE 802.3-2002)
- IGMP** Internet Group Management Protocol. A protocol through which hosts can register with their local router for multicast services. If there is more than one multicast switch/router on a given subnetwork, one of the devices is made the “querier” and assumes responsibility for keeping track of group membership.
- IGMP Proxy** Proxies multicast group membership information onto the upstream interface based on IGMP messages monitored on downstream interfaces, and forwards multicast traffic based on that information. There is no need for multicast routing protocols in a simple tree that uses IGMP Proxy.
- IGMP Query** On each subnetwork, one IGMP-capable device will act as the querier — that is, the device that asks all hosts to report on the IP multicast groups they wish to join or to which they already belong. The elected querier will be the device with the lowest IP address in the subnetwork.

IGMP Snooping Listening to IGMP Query and IGMP Report packets transferred between IP Multicast Routers and IP Multicast host groups to identify IP Multicast group members.

In-Band Management Management of the network from a station attached directly to the network.

IP Multicast Filtering A process whereby this switch can pass multicast traffic along to participating hosts.

IP Precedence The Type of Service (ToS) octet in the IPv4 header includes three precedence bits defining eight different priority levels ranging from highest priority for network control packets to lowest priority for routine traffic. The eight values are mapped one-to-one to the Class of Service categories by default, but may be configured differently to suit the requirements for specific network applications.

LACP Link Aggregation Control Protocol. Allows ports to automatically negotiate a trunked link with LACP-configured ports on another device.

Layer 2 Data Link layer in the ISO 7-Layer Data Communications Protocol. This is related directly to the hardware interface for network devices and passes on traffic based on MAC addresses.

Layer 3 Network layer in the ISO 7-Layer Data Communications Protocol. This layer handles the routing functions for data moving from one open system to another.

Link Aggregation *See Port Trunk.*

LLDP Link Layer Discovery Protocol is used to discover basic information about neighboring devices in the local broadcast domain by using periodic broadcasts to advertise information such as device identification, capabilities and configuration settings.

MD5 MD5 Message-Digest is an algorithm that is used to create digital signatures. It is intended for use with 32 bit machines and is safer than the MD4 algorithm, which has been broken. MD5 is a one-way hash function, meaning that it takes a message and converts it into a fixed string of digits, also called a message digest.

MIB Management Information Base. An acronym for Management Information Base. It is a set of database objects that contains information about a specific device.

MRD Multicast Router Discovery is a protocol used by IGMP snooping and multicast routing devices to discover which interfaces are attached to multicast routers. This process allows IGMP-enabled devices to determine where to send multicast source and group membership messages.

MSTP Multiple Spanning Tree Protocol can provide an independent spanning tree for different VLANs. It simplifies network management, provides for even faster convergence than RSTP by limiting the size of each region, and prevents VLAN members from being segmented from the rest of the group.

Multicast Switching A process whereby the switch filters incoming multicast frames for services for which no attached host has registered, or forwards them to all ports contained within the designated multicast VLAN group.

MVR Multicast VLAN Registration is a method of using a single network-wide multicast VLAN to transmit common services, such as television channels or video-on-demand, across a service-provider's network. MVR simplifies the configuration of multicast services by using a common VLAN for distribution, while still preserving security and data isolation for subscribers residing in both the MVR VLAN and other standard VLAN groups.

NTP Network Time Protocol provides the mechanisms to synchronize time across the network. The time servers operate in a hierarchical-master-slave configuration in order to synchronize local clocks within the subnet and to national time standards via wire or radio.

OSPF Open Shortest Path First is a link-state routing protocol that functions better over a larger network such as the Internet, as opposed to distance-vector routing protocols such as RIP. It includes features such as unlimited hop count, authentication of routing updates, and Variable Length Subnet Masks (VLSM).

Out-of-Band Management Management of the network from a station not attached to the network.

Port Authentication See *IEEE 802.1X*.

Port Mirroring A method whereby data on a target port is mirrored to a monitor port for troubleshooting with a logic analyzer or RMON probe. This allows data on the target port to be studied unobstructively.

Port Trunk Defines a network link aggregation and trunking method which specifies how to create a single high-speed logical link that combines several lower-speed physical links.

QinQ QinQ tunneling is designed for service providers carrying traffic for multiple customers across their networks. It is used to maintain customer-specific VLAN and Layer 2 protocol configurations even when different customers use the same internal VLAN IDs.

QoS Quality of Service. QoS refers to the capability of a network to provide better service to selected traffic flows using features such as data prioritization, queuing, congestion avoidance and traffic shaping. These features effectively provide preferential treatment to specific flows either by raising the priority of one flow or limiting the priority of another flow.

RADIUS Remote Authentication Dial-in User Service. RADIUS is a logon authentication protocol that uses software running on a central server to control access to RADIUS-compliant devices on the network.

RIP Routing Information Protocol seeks to find the shortest route to another device by minimizing the distance-vector, or hop count, which serves as a rough estimate of transmission cost. RIP-2 is a compatible upgrade to RIP. It adds useful capabilities for subnet routing, authentication, and multicast transmissions.

RMON Remote Monitoring. RMON provides comprehensive network monitoring capabilities. It eliminates the polling required in standard SNMP, and can set alarms on a variety of traffic conditions, including specific error types.

RSTP Rapid Spanning Tree Protocol. RSTP reduces the convergence time for network topology changes to about 10% of that required by the older IEEE 802.1D STP standard.

SMTP Simple Mail Transfer Protocol is a standard host-to-host mail transport protocol that operates over TCP, port 25.

SNMP Simple Network Management Protocol. The application protocol in the Internet suite of protocols which offers network management services.

SNTP Simple Network Time Protocol allows a device to set its internal clock based on periodic updates from a Network Time Protocol (NTP) server. Updates can be requested from a specific NTP server, or can be received via broadcasts sent by NTP servers.

SSH Secure Shell is a secure replacement for remote access functions, including Telnet. SSH can authenticate users with a cryptographic key, and encrypt data connections between management clients and the switch.

- STA** Spanning Tree Algorithm is a technology that checks your network for any loops. A loop can often occur in complicated or backup linked network systems. Spanning Tree detects and directs data along the shortest available path, maximizing the performance and efficiency of the network.
- TACACS+** Terminal Access Controller Access Control System Plus. TACACS+ is a logon authentication protocol that uses software running on a central server to control access to TACACS-compliant devices on the network.
- TCP/IP** Transmission Control Protocol/Internet Protocol. Protocol suite that includes TCP as the primary transport protocol, and IP as the network layer protocol.
- Telnet** Defines a remote communication facility for interfacing to a terminal device over TCP/IP.
- TFTP** Trivial File Transfer Protocol. A TCP/IP protocol commonly used for software downloads.
- UDP** User Datagram Protocol. UDP provides a datagram mode for packet-switched communications. It uses IP as the underlying transport mechanism to provide access to IP-like services. UDP packets are delivered just like IP packets – connection-less datagrams that may be discarded before reaching their targets. UDP is useful when TCP would be too complex, too slow, or just unnecessary.
- UTC** Universal Time Coordinate. UTC is a time scale that couples Greenwich Mean Time (based solely on the Earth's rotation rate) with highly accurate atomic time. The UTC does not have daylight saving time.
- VLAN** Virtual LAN. A Virtual LAN is a collection of network nodes that share the same collision domain regardless of their physical location or connection point in the network. A VLAN serves as a logical workgroup with no physical barriers, and allows users to share information and resources as though located on the same LAN.
- VRRP** Virtual Router Redundancy Protocol uses a virtual IP address to support a primary router and multiple backup routers. The backups can be configured to take over the workload if the master fails or to load share the traffic. The primary goal of VRRP is to allow a host device which has been configured with a fixed gateway to maintain network connectivity in case the primary gateway goes down.
- XModem** A protocol used to transfer files between devices. Data is grouped in 128-byte blocks and error-corrected.

Index

Numerics

(not ready)

- 802.1Q tunnel 154
 - access 161
 - configuration, guidelines 157
 - configuration, limitations 157
 - CVID to SVID map 159
 - description 154
 - ethernet type 158
 - interface configuration 161
 - mode selection 161
 - status, configuring 158
 - TPID 158
 - uplink 161
- 802.1X
 - authenticator, configuring 301
 - global settings 300
 - port authentication 298

A

- AAA
 - authorization & accounting 236
- acceptable frame type 150
- ACL 267
 - ARP 270, 280
 - binding to a port 282
 - IPv4 Extended 270, 273
 - IPv4 Standard 270, 272
 - IPv6 Extended 270, 277
 - IPv6 Standard 270, 275
 - MAC 270, 278
- Address Resolution Protocol *See* ARP
- address table 163
 - aging time 167
 - aging time, displaying 167
 - aging time, setting 167
- ARP
 - ACL 280
 - configuration 522
 - description 522
 - proxy 522
 - statistics 525

- ARP inspection 286
 - ACL filter 289
 - additional validation criteria 288
 - ARP ACL 290
 - enabling globally 288
 - trusted ports 290
- authentication
 - MAC address authentication 246
 - MAC, configuring ports 249
 - network access 246
 - public key 260
 - web 243
 - web authentication for ports, configuring 245
 - web authentication port information, displaying 245
 - web authentication, re-authenticating address 245
 - web authentication, re-authenticating ports 245
 - web, configuring 244

B

- BOOTP 482
- bootstrap router
 - PIM-SM 615
 - PIMv6-SM 631
- BPDU 174
 - filter 185
 - guard 185
 - ignoring superior BPDUs 184
 - selecting protocol based on message format 185
 - shut down port on receipt 185
- bridge extension capabilities, displaying 75
- broadcast packets, blocking 196
- broadcast storm, threshold 196, 197

C

- canonical format indicator 210
- CFM
 - basic operations 393
 - continuity check errors 425
 - continuity check messages 390, 393, 394
 - cross-check message 391, 394
 - cross-check start delay 395
 - delay measure request 414
 - description 390
 - domain service access point 391, 403, 407
 - fault isolation 390

Index

- fault notification 391, 425
- fault notification generator 393, 399, 425
- fault verification 390
- link trace cache 423
- link trace message 391, 393, 411
- loop back messages 390, 393, 412
- maintenance association 390, 403
- maintenance domain 390, 392, 398
- maintenance end point 391, 393, 398, 403, 407, 416, 417
- maintenance intermediate point 391, 399, 419
- maintenance level 391, 393
- maintenance point 390
- MEP archive hold time 401
- MEP direction 407
- remote maintenance end point 394, 403, 409, 417, 420, 421
- service instance 391, 393
- class map, DiffServ 220
- Class of Service *See* CoS
- committed burst size, QoS policy 228, 229
- committed information rate, QoS policy 228, 229
- community string 366
- configuration files, restoring defaults 77
- configuration settings
 - restoring 79, 80
 - saving 79
- Connectivity Fault Management *See* CFM
- continuity check errors, CFM 425
- continuity check messages, CFM 390, 393, 394
- CoS 199, 207
 - configuring 199
 - default mapping to internal values 210
 - enabling 206
 - layer 3/4 priorities 206
 - priorities, mapping to internal values 210
 - queue mapping 203
 - queue mode 200
 - queue weights, assigning 201
- CoS/CFI to PHB/drop precedence 210
- CPU
 - status 97
 - utilization, showing 97
- cross-check message, CFM 391, 394
- cross-check start delay, CFM 395
- CVLAN to SPVLAN map 159

D

- default IPv6 gateway, configuration 485
- default priority, ingress port 199
- default settings, system 44
- delay measure request, CFM 414
- designated router

- PIM 609
- PIMv6 625
- DHCP 482, 511
 - class identifier 511
 - client 482
 - client identifier 511
 - relay service 513
 - relay service, enabling 514
- DHCP snooping 319
 - enabling 321
 - global configuration 321
 - information option 322
 - information option policy 322
 - information option, circuit ID 325
 - information option, enabling 322
 - information option, remote ID 322
 - information option, suboption format 322
 - policy selection 322
 - specifying trusted interfaces 324
 - trusted port 324
 - untrusted port 324
 - verifying MAC addresses 321
 - VLAN configuration 323
- Differentiated Code Point Service *See* DSCP
- Differentiated Services *See* DiffServ
- DiffServ 219
 - binding policy to interface 233
 - class map 220
 - classifying QoS traffic 220
 - color aware, srTCM 228
 - color aware, trTCM 229
 - color blind, srTCM 228
 - color blind, trTCM 229
 - committed burst size 228, 229, 230
 - committed information rate 228, 229, 230
 - configuring 219
 - conforming traffic, configuring response 227
 - excess burst size 229
 - metering, configuring 224, 225
 - peak burst size 230
 - peak information rate 230
 - policy map 224
 - policy map, description 221, 227
 - QoS policy 224
 - service policy 233
 - setting CoS for matching packets 227
 - setting PHB for matching packets 227
 - single-rate, three-color meter 224, 228
 - srTCM metering 224, 228
 - traffic between CIR and BE, configuring response 228
 - traffic between CIR and PIR, configuring response 229
 - trTCM metering 229
 - two-rate, three-color meter 225
 - violating traffic, configuring response 230

DNS

- default domain name 505
- displaying the cache 510
- domain name list 505
- enabling lookup 505
- name server list 505
- static entries, IPv4 509
- static entries, IPv6 509
- Domain Name Service *See* DNS
- domain service access point, CFM 391, 403, 407
- downloading software 77
 - automatically 81
 - using FTP or TFTP 81
- DR priority, PIM-SM 611
- DR priority, PIMv6-SM 627
- drop precedence
 - CoS priority mapping 210
 - DSCP ingress map 208
- DSA encryption 262, 264, 265
- DSCP 206, 207
 - enabling 206
 - ingress map, drop precedence 208
 - mapping to internal values 207
- DSCP to PHB/drop precedence 208
- dynamic addresses
 - clearing 169
 - displaying 168
- dynamic QoS assignment 247, 250
- dynamic VLAN assignment 246, 250

E

- ECMP, maximum paths 529
- edge port, STA 184, 187
- encryption
 - DSA 262, 264, 265
 - RSA 262, 264, 265
- engine ID 356, 357
- event logging 327
- excess burst size, QoS policy 228

F

- fault isolation, CFM 390
- fault notification generator, CFM 393, 399, 425
- fault notification, CFM 391, 425
- fault verification, CFM 390
- firmware
 - displaying version 73
 - upgrading 77
 - upgrading automatically 81
 - upgrading with FTP or TFP 81
 - version, displaying 73

G

- gateway, IPv6 default 485
- general security measures 235
- GNU license 651

H

- hardware version, displaying 73
- hash mask length, PIM-SM BSR 616
- hash mask length, PIMv6-SM BSR 632
- hello holdtime
 - PIM 609
 - PIMv6 625
- hello interval
 - PIM 609
 - PIMv6 625
- HTTPS 255, 257
 - configuring 255
 - replacing SSL certificate 257
 - secure-site certificate 257
 - UDP port, configuring 256
- HTTPS, secure server 255

I

- IEEE 802.1D 173
- IEEE 802.1s 173
- IEEE 802.1w 173
- IEEE 802.1X 298
- IGMP
 - enabling per interface 474
 - filter profiles, binding to interface 460
 - filter profiles, configuration 458
 - filter, interface configuration 460
 - filter, parameters 458, 460
 - filtering & throttling 457
 - filtering & throttling, enabling 457
 - filtering & throttling, interface configuration 460
 - filtering & throttling, status 457
 - filtering, configuring profile 458
 - filtering, creating profile 458
 - filtering, group range 458
 - groups, displaying 444, 478
 - interface configuration 474
 - interface status, displaying 476
 - last member query interval 475
 - Layer 2 435
 - Layer 3 470
 - maximum response time 475
 - multicast groups, displaying 478
 - proxy 471
 - proxy routing 470
 - proxy routing, configuring 471
 - proxy routing, interface configuration 474

Index

- query 435, 437
 - query interval 475
 - query, enabling 440
 - report delay 475
 - robustness value 474
 - robustness variable 474
 - services, displaying 452, 478
 - showing groups 478
 - snooping 435
 - snooping & query, parameters 437
 - snooping, configuring 437
 - snooping, enabling 437
 - snooping, immediate leave 447
 - static groups, configuring 476
 - version 474
 - IGMP proxy
 - configuration steps 472
 - enabling 473
 - unsolicited report interval 473
 - IGMP snooping
 - configuring 445
 - enabling per interface 445, 446
 - forwarding entries 452
 - immediate leave, status 447
 - interface attached to multicast router 443, 445
 - last leave 436
 - last member query count 449
 - last member query interval 449
 - proxy query address 449
 - proxy query interval 448
 - proxy query response interval 449
 - proxy reporting 438, 448
 - querier timeout 440
 - querier, enabling 440
 - query suppression 436
 - router port expire time 440
 - static host interface 436
 - static multicast routing 441
 - static port assignment 443
 - static router interface 436
 - static router port, configuring 441
 - statistics, displaying 453
 - TCN flood 438
 - unregistered data flooding 439
 - version exclusive 439
 - version for interface, setting 448
 - version, setting 440, 448
 - with proxy reporting 436
 - immediate leave, IGMP snooping 447
 - immediate leave, MLD snooping 464
 - importing user public keys 265
 - ingress filtering 150
 - IP address
 - BOOTP/DHCP 482
 - setting 481
 - IP filter, for management access 294
 - IP Port to PHB/drop precedence 216
 - IP Precedence 207
 - enabling 206
 - IP precedence to PHB/drop precedence 214
 - IP routing 515, 543
 - configuring interfaces 518
 - maximum paths 529
 - unicast protocols 518
 - IP source guard
 - ACL table, learning mode 308
 - learning mode, ACL table or MAC table 308
 - MAC table, learning mode 308
 - IP statistics 496
 - IPv4 address
 - BOOTP/DHCP 482
 - setting 481
 - IPv4 source guard
 - configuring static entries 309
 - setting filter criteria 307
 - setting maximum bindings 308
 - IPv6
 - displaying neighbors 495
 - duplicate address detection 495
 - enabling 487
 - MTU 487
 - neighbor solicitation interval 488
 - reachability time 489
 - router advertisements, blocking 489
 - IPv6 address
 - dynamic configuration (link-local) 487
 - EUI format 492
 - EUI-64 setting 492
 - explicit configuration 487
 - global unicast 492
 - link-local 493
 - manual configuration (global unicast) 492
 - manual configuration (link-local) 493
 - setting 485
 - IPv6 source guard
 - configuring static entries 316
 - setting filter criteria 313
 - setting maximum bindings 315
- ## J
- jumbo frame 74
- ## K
- key
 - private 258
 - public 258
 - user public, importing 265

- key pair
 - host 258
 - host, generating 262, 264

L

LACP

- admin key 130
- configuration 129
- group attributes, configuring 133
- group members, configuring 131
- load balancing 139
- local parameters 136
- partner parameters 138
- protocol message statistics 135
- protocol parameters 129
- timeout mode 130
- timeout, for LACPDU 130

last member query count, IGMP snooping 449

last member query interval, IGMP snooping 449

license information, GNU 651

Link Layer Discovery Protocol - Media Endpoint Discovery
See LLDP-MED

Link Layer Discovery Protocol See LLDP

link trace cache, CFM 423

link trace message, CFM 391, 393, 411

link type, STA 183, 187

LLDP 331

- device statistics details, displaying 353
- device statistics, displaying 351
- display device information 339, 343
- displaying remote information 343
- interface attributes, configuring 334
- local device information, displaying 339
- message attributes 334
- message statistics 351
- remote information, displaying 350, 351
- remote port information, displaying 343
- timing attributes, configuring 332
- TLV 331, 334
- TLV, 802.1 335
- TLV, 802.3 335
- TLV, basic 334
- TLV, management address 334
- TLV, port description 335
- TLV, system capabilities 335
- TLV, system description 335
- TLV, system name 335

LLDP-MED 331

- end-node, extended power-via-MDI 348
- end-node, inventory 348
- end-node, location 348
- end-node, network policy 347
- notification, status 334
- TLV 336

- TLV, civic address 336, 337

- TLV, inventory 336

- TLV, location 336

- TLV, MED capabilities 336

- TLV, network policy 336

local engine ID 356

logging

- messages, displaying 329

- syslog traps 330

- to syslog servers 330

log-in, web interface 50

logon authentication 241

- encryption keys 240

- RADIUS client 239

- RADIUS server 239

- sequence 237

- settings 238

- TACACS+ client 238

- TACACS+ server 238

logon authentication, settings 239

loop back messages, CFM 390, 393, 412

M

MAC address authentication 246

- ports, configuring 249

- reauthentication 249

main menu, web interface 52

maintenance association, CFM 390, 403

maintenance domain, CFM 390, 392, 398

maintenance end point, CFM 391, 393, 398, 403, 407, 416, 417

maintenance intermediate point, CFM 391, 399, 419

maintenance level, CFM 391, 393

maintenance point, CFM 390

management access, filtering per address 294

management access, IP filter 294

Management Information Bases (MIBs) 646

matching class settings, classifying QoS traffic 221

media-type 105, 107

memory

- status 98

- utilization, showing 98

MEP archive hold time, CFM 401

mirror port

- configuring local traffic 108

- configuring remote traffic 110

mirror port, configuring 108

MLD snooping 462

- configuring 462

- enabling 462

- groups, displaying 468, 469

- immediate leave 464

- immediate leave, status 464

- interface attached to multicast router 465, 466

Index

- multicast static router port 465
 - querier 462
 - querier, enabling 462
 - query interval 463
 - query, maximum response time 463
 - robustness value 463
 - static port assignment 467
 - static router port 465
 - unknown multicast, handling 463
 - version 463
- MSTP 173, 188
 - global settings, configuring 175, 188
 - global settings, displaying 180
 - interface settings, configuring 181, 192
 - interface settings, displaying 194
 - max hop count 178
 - path cost 193
 - region name 178
 - region revision 178
- MTU for IPv6 487
- multicast filtering 433
 - enabling IGMP snooping 437, 446
 - enabling IGMP snooping per interface 445
 - enabling MLD snooping 462
 - router configuration 441
- multicast groups 444, 452, 468, 478
 - displaying 444, 452, 468, 478
 - static 443, 444, 467, 468
- multicast router discovery 446
- multicast router port, displaying 442, 466
- multicast routing 599
 - description 599
 - enabling, IPv4 602
 - enabling, IPv6 602
 - global settings, IPv4 602
 - global settings, IPv6 602
 - PIM 607
 - PIM-DM 607
 - PIM-SM 607, 614
 - PIMv6 623
 - PIMv6-SM 630
 - reverse path tree 600
 - routing table, IPv4 603
 - upstream interface 603
- multicast services
 - configuring 443, 467
 - displaying 444, 468
- multicast static router port 441
 - configuring 441
 - configuring for MLD snooping 465
- multicast storm, threshold 196, 197
- multicast, filtering and throttling 457
- multicast, static router port 441

N

- network access
 - authentication 246
 - dynamic QoS assignment 250
 - dynamic VLAN assignment 250
 - guest VLAN 250
 - MAC address filter 250
 - port configuration 249
 - reauthentication 249
 - secure MAC information 254

NTP

- authentication keys, specifying 90
- setting the system clock 89
- specifying servers 89

O

OSPF 562

- ABR route summary 578
- area border router 563, 570, 572, 573, 574, 576, 577, 578, 590
- AS summary route 582
- authentication key 586
- authentication type 586
- auto cost for an interface 567
- autonomous system boundary router 568, 570, 572, 574, 577, 580, 582, 593
- backbone 563, 564, 572, 589, 590
- configuration settings, displaying 569, 589
- cost for interface 584
- default cost for summary route 574, 576
- default external route 568
- default metric for external routes 567
- enabling 565
- general settings 566, 569
- hello interval 585
- interface summary information, displaying 589
- LSA advertisement interval 586
- LSA database, displaying 592
- message digest key 586
- neighboring router information, displaying 595
- network area 563
- normal area 564
- NSSA 571, 572, 577
- process ID 565, 566, 571, 573, 576, 577, 579, 581, 583
- process parameters, displaying 569
- redistributing external routes 580
- retransmit interval 586
- RFC 1583 compatible 566
- route summary, ABR 578
- router ID 567
- router priority 585
- routing table, displaying 592

- SPF timers 567
- stub 571, 575
- transit area 563, 564, 573, 575, 589, 590
- transmit delay over interface 585
- virtual link 589
- virtual links, displaying 591

P

- packet block
 - broadcast 197
 - multicast 196
 - unknown multicast 197
 - unknown unicast 196, 197
- passwords 241
 - administrator setting 241
- path cost 187
 - method 177
 - STA 182, 187
- peak burst size, QoS policy 229
- peak information rate, QoS policy 229
- per-hop behavior, DSCP ingress map 208
- PIM 607
 - configuring 607
 - dense-mode attributes 611
 - designated router 609
 - enabling for interfaces 608, 609
 - enabling globally 607
 - hello holdtime 609
 - hello interval 609
 - interface settings 608
 - neighbor routers, displaying 613
 - sparse-mode attributes 611
- PIM-DM 607
 - configuring 607
 - global configuration 609–611, 614
 - interface settings 611
 - neighbor routers 613
- PIM-SM 607, 614
 - bootstrap router 615
 - BSR candidate 615
 - BSR elected, displaying 620
 - configuring 607, 614
 - DR priority 611
 - global configuration 611
 - hash mask length for BSR 616
 - interface settings 611
 - neighbor routers 613
 - register rate limit for DR 614
 - rendezvous point 617
 - RP candidate 618
 - RP candidate, advertising 618
 - RP mapping, displaying 622
 - shared tree 614
 - shortest path tree 614
 - SPT threshold 614
 - static RP, configuring 617
- PIMv6 623
 - configuring 623
 - dense mode, enabling 625
 - dense-mode attributes 627
 - designated router 625
 - enabling for interfaces 624
 - enabling globally 623
 - global configuration 623
 - graft retry interval 627
 - hello holdtime 625
 - hello interval 625
 - interface configuration, displaying 628, 629
 - interface settings 624
 - max graft retries 627
 - neighbor routers 629
 - neighbor routers, displaying 629
 - prune delay 626
 - prune state, hold time 625
 - sparse-mode attributes 627
 - state refresh message interval 627
 - triggered hello delay 626
- PIMv6-DM
 - global configuration 625, 630
 - interface settings 627
- PIMv6-SM 630
 - bootstrap router 631
 - BSR candidate 631
 - BSR elected, displaying 637
 - configuring 630
 - DR priority 627
 - global configuration 627
 - hash mask length for BSR 632
 - interface settings 627
 - register rate limit for DR 630
 - rendezvous point 633
 - RP candidate 635
 - RP candidate, advertising 635
 - RP mapping, displaying 638
 - shared tree 631
 - shortest path tree 630
 - SPT threshold 630
 - static RP, configuring 633
- policing traffic, QoS policy 224, 228
- policy map
 - description 227
 - DiffServ 224
- port authentication 298
- port priority
 - configuring 199
 - default ingress 199
 - STA 182
- port security, configuring 296

Index

ports

- autonegotiation 105
 - broadcast storm threshold 196, 197
 - capabilities 105
 - configuring 104
 - duplex mode 105
 - flow control 105
 - forced selection of media type 105, 107
 - mirroring 108
 - mirroring local traffic 108
 - mirroring remote traffic 110
 - mtu 105
 - multicast storm threshold 196, 197
 - speed 105
 - statistics 114
 - unknown unicast storm threshold 196, 197
- priority, default port ingress 199
- private key 258
- problems, troubleshooting 649
- protocol migration 185
- proxy ARP 522
- proxy query address, IGMP snooping 449
- proxy query interval, IGMP snooping 448
- proxy query response interval, IGMP snooping 449
- proxy reporting, IGMP snooping 448
- public key 258
- PVID, port native VLAN 150

Q

QinQ Tunneling *See* 802.1Q tunnel

QoS 219

- configuration guidelines 220
 - configuring 219
 - CoS/CFI to PHB/drop precedence 210
 - DSCP to PHB/drop precedence 207
 - dynamic assignment 250
 - IP Port to PHB/drop precedence 216
 - IP precedence to PHB/drop precedence 214
 - matching class settings 221
 - PHB to queue 203
 - selecting CoS, DSCP, IP Precedence 206
- QoS policy
- committed burst size 228, 229
 - committed information rate 228, 229
 - excess burst size 228
 - peak burst size 229
 - peak information rate 229
 - policing flow 224, 228
 - srTCM 224
 - srTCM police meter 228
 - trTCM 225
 - trTCM police meter 229
- QoS policy, committed information rate 228
- Quality of Service *See* QoS

- queue mode, setting 200
- queue weight, assigning to CoS 201

R

RADIUS

- logon authentication 239
 - settings 239
- rate limit
- port 195
 - setting 195
- register rate limit, PIM-SM 614
- register rate limit, PIMv6-SM 630
- remote engine ID 357
- remote logging 330
- remote maintenance end point, CFM 394, 403, 409, 417, 420, 421
- rendezvous point
- PIM-SM 617
 - PIMv6-SM 633
- restarting the system 98
- at scheduled times 98
- RIP 544
- authentication key 558
 - authentication mode 558
 - clearing routes 548
 - configuring 544
 - default external route 546
 - default metric 546
 - description 543
 - global settings 545
 - interface protocol settings 556
 - interface, enabling 549
 - neighbor router 552, 561
 - passively monitoring updates 551, 596
 - poison reverse 544, 559
 - protocol packets, receiving 558
 - protocol packets, sending 558
 - receive version 558
 - redistributing external routing information 553
 - routes, clearing 548
 - routes, displaying 561
 - routing table, clearing 548
 - send version 558
 - specifying interfaces 549
 - split horizon 544, 559
 - timers 547
 - version 545
- RMON 380
- alarm, displaying settings 382
 - alarm, setting thresholds 380
 - event settings, displaying 385
 - response to alarm setting 383
 - statistics history, collection 385
 - statistics history, displaying 387

- statistics, collection 388
 - statistics, displaying 389
 - router redundancy
 - protocols 533
 - VRRP 533
 - routing table, displaying 528
 - RSA encryption 262, 264, 265
 - RSTP 173
 - global settings, configuring 175
 - global settings, displaying 180
 - interface settings, configuring 181
 - interface settings, displaying 186
- S**
- secure shell 258
 - configuration 258
 - security, general measures 235
 - serial port, configuring 93
 - service instance, CFM 391, 393
 - shared tree
 - PIM-SM 614
 - PIMv6-SM 631
 - shortest path tree
 - PIM-SM 614
 - PIMv6-SM 630
 - Simple Network Management Protocol *See* SNMP
 - single rate three color meter *See* srTCM
 - SNMP 353
 - community string 366
 - enabling traps 372
 - enabling traps, mac-address changes 170
 - filtering IP addresses 294
 - global settings, configuring 355
 - trap manager 372
 - users, configuring 367, 369
 - SNMPv3
 - engine ID 356, 357
 - engine identifier, local 356
 - engine identifier, remote 357
 - groups 361
 - local users, configuring 367
 - remote users, configuring 369
 - user configuration 367, 369
 - views 358
 - SNTP
 - setting the system clock 86
 - specifying servers 88
 - software
 - displaying version 73
 - downloading 77
 - version, displaying 73
 - Spanning Tree Protocol *See* STA
 - specifications, software 643
 - SPT threshold, PIM-SM 614
 - SPT threshold, PIMv6-SM 630
 - srTCM
 - police meter 228
 - QoS policy 224
 - SSH 258
 - authentication retries 261
 - configuring 258
 - downloading public keys for clients 265
 - generating host key pair 262, 264
 - server, configuring 261
 - timeout 261
 - SSL, replacing certificate 257
 - STA 173
 - BPDU filter 185
 - BPDU flooding 182
 - BPDU shutdown 185
 - edge port 184, 187
 - forward delay 178
 - global settings, configuring 175
 - global settings, displaying 180
 - hello time 177
 - interface settings, configuring 181
 - interface settings, displaying 186
 - link type 183, 187
 - maximum age 177
 - MSTP interface settings, configuring 192
 - MSTP path cost 193
 - path cost 182, 187
 - path cost method 177
 - port priority 182
 - protocol migration 185
 - transmission limit 177
 - standards, IEEE 645
 - startup files
 - creating 77
 - displaying 77
 - setting 77
 - static addresses, setting 165
 - static routes, configuring 526
 - statistics
 - ARP 525
 - history for port 118
 - history for trunk 118
 - port 114
 - STP 176
 - switch settings
 - restoring 79
 - saving 79
 - system clock
 - setting 85
 - setting manually 85
 - setting the time zone 92
 - setting with NTP 89
 - setting with SNTP 86

Index

system logs 327
system software, downloading from server 77

T

TACACS+

logon authentication 238
settings 239

TCN

flood 438
general query solicitation 439

Telnet

configuring 95
server, enabling 95

telnet connection, configuring 95

time zone, setting 92

time, setting 85

TPID 158

traffic segmentation 141

assigning ports 141
enabling 141
sessions, assigning ports 143
sessions, creating 142

transceiver data

configuring trap thresholds 124
displaying 123

transceiver data, displaying 122

trap manager 372

troubleshooting 649

trTCM

police meter 229
QoS policy 225

trunk

configuration 125
LACP 129
static 126

two rate three color meter *See* trTCM

Type Length Value *See* LLDP TLV

U

UDLD

configuration 427
interface settings 429
neighbor information 430
protocol intervals 427

unicast routing 543

ECMP 529
ECMP maximum paths 529

unidirectional link detection 427

unknown multicast packets, blocking 196

unknown unicast packets, blocking 196
unknown unicast storm, threshold 196, 197
unregistered data flooding, IGMP snooping 439
upgrading software 77

upstream interface, multicast route 603

user account 241

user password 241

V

VLANs 145–147

802.1Q tunnel mode 161
acceptable frame type 150
adding static members 150
configuring port members, VLAN index 152
creating 147
description 145, 147
displaying port members 152, 153
displaying port members by interface 153
displaying port members by interface range 153
displaying port members by VLAN index 152
dynamic assignment 250
egress mode 150
ingress filtering 150
interface configuration 150
port members, displaying 152, 153
PVID 150

VRRP 533

authentication 537
configuration settings 533, 534
group statistics 541
preemption 535, 537
priority 535, 537
protocol message statistics 540
timers 536
virtual address 535

W

web authentication 243

address, re-authenticating 245
configuring 244
port information, displaying 245
ports, configuring 245
ports, re-authenticating 245

web interface

access requirements 49
configuration buttons 51
home page 50
menu list 52
panel display 51

