



GTL-2891

**28-Port Stackable Layer 3 Managed Gigabit Switch,
2 x SFP+, 2 x SFP+ (Optional Modules)**

User Manual

V1.0

Digital Data Communications Asia Co., Ltd.

<http://www.level1.com>

User Manual

GTL-2891

Layer 3 Stackable Gigabit Ethernet Switch
with 24 10/100/1000BASE-T (RJ-45) Ports,
2 10-Gigabit SFP+ Ports, and
Optional Module with 2 10-Gigabit SFP+ Ports

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 provides detailed information about the switch's key features. It also describes the switch's web browser interface. For information on the command line interface refer to the *CLI Reference 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.

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

October 2016 Revision

This is the first version of this guide. This guide is valid for software release v1.2.2.19a.

Contents

How to Use This Guide	3
Contents	5
Figures	19
Tables	35

Section I	Getting Started	37
	1 Introduction	39
	Key Features	39
	Description of Software Features	41
	IP Routing	45
	Equal-cost Multipath Load Balancing	46
	Router Redundancy	46
	Address Resolution Protocol	46
	Operation, Administration, and Maintenance	46
	System Defaults	47

Section II	Web Configuration	51
	2 Using the Web Interface	53
	Connecting to the Web Interface	53
	Navigating the Web Browser Interface	54
	Home Page	54
	Configuration Options	55
	Panel Display	55
	Main Menu	56

3 Basic Management Tasks	79
Displaying System Information	80
Displaying Hardware/Software Versions	81
Configuring Support for Jumbo Frames	82
Displaying Bridge Extension Capabilities	83
Managing System Files	85
Copying Files via FTP/TFTP or HTTP	85
Saving the Running Configuration to a Local File	87
Setting the Start-Up File	88
Showing System Files	89
Automatic Operation Code Upgrade	89
Setting the System Clock	93
Setting the Time Manually	94
Setting the SNTP Polling Interval	95
Configuring NTP	95
Configuring Time Servers	96
Setting the Time Zone	100
Configuring Summer Time	101
Configuring the Console Port	103
Configuring Telnet Settings	105
Displaying CPU Utilization	106
Displaying Memory Utilization	107
Stacking	108
Setting the Master Unit	108
Enabling Stacking Ports	109
Renumbering the Stack	110
Resetting the System	111
4 Interface Configuration	115
Port Configuration	116
Configuring by Port List	116
Configuring by Port Range	118
Displaying Connection Status	119
Configuring Local Port Mirroring	120
Configuring Remote Port Mirroring	122

Showing Port or Trunk Statistics	126
Displaying Transceiver Data	131
Configuring Transceiver Thresholds	132
Performing Cable Diagnostics	134
Trunk Configuration	135
Configuring a Static Trunk	137
Configuring a Dynamic Trunk	139
Displaying LACP Port Counters	145
Displaying LACP Settings and Status for the Local Side	146
Displaying LACP Settings and Status for the Remote Side	148
Configuring Load Balancing	149
Saving Power	151
Traffic Segmentation	153
Enabling Traffic Segmentation	153
Configuring Uplink and Downlink Ports	154
VLAN Trunking	156
5 VLAN Configuration	159
IEEE 802.1Q VLANs	159
Configuring VLAN Groups	162
Adding Static Members to VLANs	165
Configuring Dynamic VLAN Registration	169
IEEE 802.1Q Tunneling	172
Enabling QinQ Tunneling on the Switch	176
Creating CVLAN to SPVLAN Mapping Entries	178
Adding an Interface to a QinQ Tunnel	179
Protocol VLANs	181
Configuring Protocol VLAN Groups	181
Mapping Protocol Groups to Interfaces	183
Configuring IP Subnet VLANs	185
Configuring MAC-based VLANs	187
Configuring VLAN Mirroring	189
Configuring VLAN Translation	191
6 Address Table Settings	193
Configuring MAC Address Learning	193

Setting Static Addresses	195
Changing the Aging Time	197
Displaying the Dynamic Address Table	197
Clearing the Dynamic Address Table	199
Configuring MAC Address Mirroring	200
Issuing MAC Address Traps	201
7 Spanning Tree Algorithm	203
Overview	203
Configuring Loopback Detection	205
Configuring Global Settings for STA	207
Displaying Global Settings for STA	212
Configuring Interface Settings for STA	213
Displaying Interface Settings for STA	218
Configuring Multiple Spanning Trees	220
Configuring Interface Settings for MSTP	225
8 Congestion Control	227
Rate Limiting	227
Storm Control	228
Automatic Traffic Control	230
Setting the ATC Timers	231
Configuring ATC Thresholds and Responses	233
9 Class of Service	237
Layer 2 Queue Settings	237
Setting the Default Priority for Interfaces	237
Selecting the Queue Mode	238
Mapping CoS Values to Egress Queues	241
Layer 3/4 Priority Settings	244
Setting Priority Processing to DSCP or CoS	244
Mapping Ingress DSCP Values to Internal DSCP Values	245
Mapping CoS Priorities to Internal DSCP Values	248
10 Quality of Service	251
Overview	251
Configuring a Class Map	252

Creating QoS Policies	256
Attaching a Policy Map to a Port	265
11 VoIP Traffic Configuration	267
Overview	267
Configuring VoIP Traffic	268
Configuring Telephony OUI	269
Configuring VoIP Traffic Ports	270
12 Security Measures	273
AAA (Authentication, Authorization and Accounting)	274
Configuring Local/Remote Logon Authentication	275
Configuring Remote Logon Authentication Servers	276
Configuring AAA Accounting	281
Configuring AAA Authorization	287
Configuring User Accounts	290
Web Authentication	292
Configuring Global Settings for Web Authentication	293
Configuring Interface Settings for Web Authentication	294
Network Access (MAC Address Authentication)	295
Configuring Global Settings for Network Access	297
Configuring Network Access for Ports	298
Configuring Port Link Detection	300
Configuring a MAC Address Filter	301
Displaying Secure MAC Address Information	303
Configuring HTTPS	305
Configuring Global Settings for HTTPS	305
Replacing the Default Secure-site Certificate	306
Configuring the Secure Shell	308
Configuring the SSH Server	310
Generating the Host Key Pair	312
Importing User Public Keys	313
Access Control Lists	315
Setting a Time Range	316
Showing TCAM Utilization	319
Setting the ACL Name and Type	320

Configuring a Standard IPv4 ACL	322
Configuring an Extended IPv4 ACL	323
Configuring a Standard IPv6 ACL	325
Configuring an Extended IPv6 ACL	327
Configuring a MAC ACL	329
Configuring an ARP ACL	331
Binding a Port to an Access Control List	332
Configuring ACL Mirroring	333
Showing ACL Hardware Counters	335
ARP Inspection	336
Configuring Global Settings for ARP Inspection	337
Configuring VLAN Settings for ARP Inspection	339
Configuring Interface Settings for ARP Inspection	341
Displaying ARP Inspection Statistics	342
Displaying the ARP Inspection Log	343
Filtering IP Addresses for Management Access	344
Configuring Port Security	346
Configuring 802.1X Port Authentication	348
Configuring 802.1X Global Settings	350
Configuring Port Authenticator Settings for 802.1X	351
Configuring Port Supplicant Settings for 802.1X	356
Displaying 802.1X Statistics	358
DoS Protection	360
IPv4 Source Guard	363
Configuring Ports for IPv4 Source Guard	363
Configuring Static Bindings for IPv4 Source Guard	365
Displaying Information for Dynamic IPv4 Source Guard Bindings	368
IPv6 Source Guard	369
Configuring Ports for IPv6 Source Guard	369
Configuring Static Bindings for IPv6 Source Guard	371
Displaying Information for Dynamic IPv6 Source Guard Bindings	374
DHCP Snooping	375
DHCP Snooping Global Configuration	377
DHCP Snooping VLAN Configuration	379
Configuring Ports for DHCP Snooping	380

Displaying DHCP Snooping Binding Information	381
13 Basic Administration Protocols	383
Configuring Event Logging	384
System Log Configuration	384
Remote Log Configuration	386
Sending Simple Mail Transfer Protocol Alerts	387
Link Layer Discovery Protocol	389
Setting LLDP Timing Attributes	389
Configuring LLDP Interface Attributes	391
Configuring LLDP Interface Civic-Address	395
Displaying LLDP Local Device Information	397
Displaying LLDP Remote Device Information	400
Displaying Device Statistics	409
Simple Network Management Protocol	411
Configuring Global Settings for SNMP	413
Setting the Local Engine ID	414
Specifying a Remote Engine ID	415
Setting SNMPv3 Views	417
Configuring SNMPv3 Groups	419
Setting Community Access Strings	425
Configuring Local SNMPv3 Users	426
Configuring Remote SNMPv3 Users	429
Specifying Trap Managers	432
Creating SNMP Notification Logs	436
Showing SNMP Statistics	438
Remote Monitoring	440
Configuring RMON Alarms	440
Configuring RMON Events	443
Configuring RMON History Samples	445
Configuring RMON Statistical Samples	448
Switch Clustering	450
Configuring General Settings for Clusters	451
Cluster Member Configuration	452
Managing Cluster Members	454

Ethernet Ring Protection Switching	455
ERPS Global Configuration	459
ERPS Ring Configuration	460
ERPS Forced and Manual Mode Operations	476
Connectivity Fault Management	480
Configuring Global Settings for CFM	484
Configuring Interfaces for CFM	487
Configuring CFM Maintenance Domains	487
Configuring CFM Maintenance Associations	492
Configuring Maintenance End Points	496
Configuring Remote Maintenance End Points	498
Transmitting Link Trace Messages	500
Transmitting Loop Back Messages	502
Transmitting Delay-Measure Requests	504
Displaying Local MEPs	506
Displaying Details for Local MEPs	507
Displaying Local MIPs	509
Displaying Remote MEPs	510
Displaying Details for Remote MEPs	511
Displaying the Link Trace Cache	513
Displaying Fault Notification Settings	514
Displaying Continuity Check Errors	515
OAM Configuration	516
Enabling OAM on Local Ports	516
Displaying Statistics for OAM Messages	519
Displaying the OAM Event Log	520
Displaying the Status of Remote Interfaces	521
Configuring a Remote Loopback Test	522
Displaying Results of Remote Loopback Testing	524
UDLD Configuration	525
Configuring UDLD Protocol Intervals	526
Configuring UDLD Interface Settings	527
Displaying UDLD Neighbor Information	529

14 Multicast Filtering	531
Overview	531
IGMP Protocol	533
Layer 2 IGMP (Snooping and Query for IPv4)	533
Configuring IGMP Snooping and Query Parameters	535
Specifying Static Interfaces for a Multicast Router	539
Assigning Interfaces to Multicast Services	541
Setting IGMP Snooping Status per Interface	544
Filtering IGMP Query Packets and Multicast Data	550
Displaying Multicast Groups Discovered by IGMP Snooping	551
Displaying IGMP Snooping Statistics	552
Filtering and Throttling IGMP Groups	556
Enabling IGMP Filtering and Throttling	556
Configuring IGMP Filter Profiles	557
Configuring IGMP Filtering and Throttling for Interfaces	559
MLD Snooping (Snooping and Query for IPv6)	561
Configuring MLD Snooping and Query Parameters	561
Setting Immediate Leave Status for MLD Snooping per Interface	563
Specifying Static Interfaces for an IPv6 Multicast Router	564
Assigning Interfaces to IPv6 Multicast Services	566
Showing MLD Snooping Groups and Source List	568
Layer 3 IGMP (Query used with Multicast Routing)	569
Configuring IGMP Proxy Routing	570
Configuring IGMP Interface Parameters	573
Configuring Static IGMP Group Membership	575
Displaying Multicast Group Information	577
Multicast VLAN Registration for IPv4	579
Configuring MVR Global Settings	581
Configuring MVR Domain Settings	583
Configuring MVR Group Address Profiles	584
Configuring MVR Interface Status	587
Assigning Static MVR Multicast Groups to Interfaces	589
Displaying MVR Receiver Groups	591
Displaying MVR Statistics	592

Multicast VLAN Registration for IPv6	596
Configuring MVR6 Global Settings	597
Configuring MVR6 Domain Settings	599
Configuring MVR6 Group Address Profiles	600
Configuring MVR6 Interface Status	603
Assigning Static MVR6 Multicast Groups to Interfaces	605
Displaying MVR6 Receiver Groups	607
Displaying MVR6 Statistics	608
15 IP Configuration	613
Setting the Switch's IP Address (IP Version 4)	613
Sending DHCP Inform Requests for Additional Information	617
Setting the Switch's IP Address (IP Version 6)	618
Configuring the IPv6 Default Gateway	619
Configuring IPv6 Interface Settings	620
Configuring an IPv6 Address	624
Showing IPv6 Addresses	626
Showing the IPv6 Neighbor Cache	628
Showing IPv6 Statistics	629
Showing the MTU for Responding Destinations	635
16 IP Services	637
Domain Name Service	637
Configuring General DNS Service Parameters	637
Configuring a List of Domain Names	638
Configuring a List of Name Servers	640
Configuring Static DNS Host to Address Entries	641
Displaying the DNS Cache	642
Dynamic Host Configuration Protocol	643
Specifying a DHCP Client Identifier	643
Configuring DHCP Relay Service	645
Configuring the DHCP Server	646
Forwarding UDP Service Requests	653
Enabling the UDP Helper	654
Specifying UDP Destination Ports	654
Specifying the Target Server or Subnet	655

Configuring the PPPoE Intermediate Agent	657
Configuring PPPoE IA Global Settings	657
Configuring PPPoE IA Interface Settings	659
Showing PPPoE IA Statistics	660
17 General IP Routing	663
Overview	663
Initial Configuration	663
IP Routing and Switching	664
Routing Path Management	665
Routing Protocols	666
Configuring IP Routing Interfaces	666
Configuring Local and Remote Interfaces	666
Using the Ping Function	667
Using the Trace Route Function	668
Address Resolution Protocol	670
Basic ARP Configuration	670
Configuring Static ARP Addresses	672
Displaying Dynamic or Local ARP Entries	674
Displaying ARP Statistics	674
Configuring Static Routes	675
Displaying the Routing Table	677
Equal-cost Multipath Routing	678
18 Configuring Router Redundancy	681
Configuring VRRP Groups	682
Displaying VRRP Global Statistics	688
Displaying VRRP Group Statistics	689
19 Unicast Routing	691
Overview	691
Configuring the Routing Information Protocol	692
Configuring General Protocol Settings	693
Clearing Entries from the Routing Table	696
Specifying Network Interfaces	697
Specifying Passive Interfaces	699

Specifying Static Neighbors	700
Configuring Route Redistribution	701
Specifying an Administrative Distance	703
Configuring Network Interfaces for RIP	704
Displaying RIP Interface Settings	708
Displaying Peer Router Information	709
Resetting RIP Statistics	709
Configuring the Open Shortest Path First Protocol (Version 2)	710
Defining Network Areas Based on Addresses	711
Configuring General Protocol Settings	714
Displaying Administrative Settings and Statistics	717
Adding an NSSA or Stub	719
Configuring NSSA Settings	720
Configuring Stub Settings	723
Displaying Information on NSSA and Stub Areas	725
Configuring Area Ranges (Route Summarization for ABRs)	726
Redistributing External Routes	728
Configuring Summary Addresses (for External AS Routes)	730
Configuring OSPF Interfaces	732
Configuring Virtual Links	737
Displaying Link State Database Information	740
Displaying Information on Neighboring Routers	743
20 Multicast Routing	745
Overview	745
Configuring Global Settings for Multicast Routing	748
Enabling Multicast Routing Globally	748
Displaying the Multicast Routing Table	749
Configuring PIM for IPv4	753
Enabling PIM Globally	753
Configuring PIM Interface Settings	754
Displaying PIM Neighbor Information	759
Configuring Global PIM-SM Settings	760
Configuring a PIM BSR Candidate	761
Configuring a PIM Static Rendezvous Point	763

Configuring a PIM RP Candidate	764
Displaying the PIM BSR Router	766
Displaying PIM RP Mapping	768
Configuring PIMv6 for IPv6	769
Enabling PIMv6 Globally	769
Configuring PIMv6 Interface Settings	769
Displaying PIM6 Neighbor Information	775
Configuring Global PIM6-SM Settings	775
Configuring a PIM6 BSR Candidate	777
Configuring a PIM6 Static Rendezvous Point	778
Configuring a PIM6 RP Candidate	780
Displaying the PIM6 BSR Router	782
Displaying RP Mapping	783

Section III	Appendices	785
A	Software Specifications	787
	Software Features	787
	Management Features	788
	Standards	789
	Management Information Bases	790
B	Troubleshooting	793
	Problems Accessing the Management Interface	793
	Using System Logs	794
C	License Statement / GPL Code Statement	795
	Written Offer for GPL/LGPL Source Code	795
	The GNU General Public License	795
	How to Apply These Terms to Your New Programs	799
	Notification of Compliance	800
	Glossary	803
	Index	811

Figures

Figure 1: Home Page	54
Figure 2: Front Panel Indicators	55
Figure 3: System Information	80
Figure 4: General Switch Information	82
Figure 5: Configuring Support for Jumbo Frames	83
Figure 6: Displaying Bridge Extension Configuration	84
Figure 7: Copy Firmware	87
Figure 8: Saving the Running Configuration	88
Figure 9: Setting Start-Up Files	88
Figure 10: Displaying System Files	89
Figure 11: Configuring Automatic Code Upgrade	93
Figure 12: Manually Setting the System Clock	94
Figure 13: Setting the Polling Interval for SNTP	95
Figure 14: Configuring NTP	96
Figure 15: Specifying SNTP Time Servers	97
Figure 16: Adding an NTP Time Server	98
Figure 17: Showing the NTP Time Server List	98
Figure 18: Adding an NTP Authentication Key	99
Figure 19: Showing the NTP Authentication Key List	100
Figure 20: Setting the Time Zone	101
Figure 21: Configuring Summer Time	103
Figure 22: Console Port Settings	104
Figure 23: Telnet Connection Settings	106
Figure 24: Displaying CPU Utilization	107
Figure 25: Displaying Memory Utilization	107
Figure 26: Setting the Stack Master	109
Figure 27: Enabling Stacking on 10G Ports	110
Figure 28: Renumbering the Stack	110
Figure 29: Restarting the Switch (Immediately)	113

Figure 30: Restarting the Switch (In)	113
Figure 31: Restarting the Switch (At)	114
Figure 32: Restarting the Switch (Regularly)	114
Figure 33: Configuring Connections by Port List	118
Figure 34: Configuring Connections by Port Range	119
Figure 35: Displaying Port Information	120
Figure 36: Configuring Local Port Mirroring	120
Figure 37: Configuring Local Port Mirroring	121
Figure 38: Displaying Local Port Mirror Sessions	122
Figure 39: Configuring Remote Port Mirroring	122
Figure 40: Configuring Remote Port Mirroring (Source)	125
Figure 41: Configuring Remote Port Mirroring (Intermediate)	126
Figure 42: Configuring Remote Port Mirroring (Destination)	126
Figure 43: Showing Port Statistics (Table)	129
Figure 44: Showing Port Statistics (Chart)	130
Figure 45: Displaying Transceiver Data	131
Figure 46: Configuring Transceiver Thresholds	133
Figure 47: Performing Cable Tests	135
Figure 48: Configuring Static Trunks	137
Figure 49: Creating Static Trunks	138
Figure 50: Adding Static Trunks Members	138
Figure 51: Configuring Connection Parameters for a Static Trunk	139
Figure 52: Showing Information for Static Trunks	139
Figure 53: Configuring Dynamic Trunks	139
Figure 54: Configuring the LACP Aggregator Admin Key	142
Figure 55: Enabling LACP on a Port	143
Figure 56: Configuring LACP Parameters on a Port	144
Figure 57: Showing Members of a Dynamic Trunk	144
Figure 58: Configuring Connection Settings for a Dynamic Trunk	145
Figure 59: Showing Connection Parameters for Dynamic Trunks	145
Figure 60: Displaying LACP Port Counters	146
Figure 61: Displaying LACP Port Internal Information	148
Figure 62: Displaying LACP Port Remote Information	149
Figure 63: Configuring Load Balancing	151
Figure 64: Enabling Power Savings	152

Figure 65: Enabling Traffic Segmentation	154
Figure 66: Configuring Members for Traffic Segmentation	155
Figure 67: Showing Traffic Segmentation Members	156
Figure 68: Configuring VLAN Trunking	156
Figure 69: Configuring VLAN Trunking	158
Figure 70: VLAN Compliant and VLAN Non-compliant Devices	161
Figure 71: Using GVRP	162
Figure 72: Creating Static VLANs	164
Figure 73: Modifying Settings for Static VLANs	164
Figure 74: Showing Static VLANs	165
Figure 75: Configuring Static Members by VLAN Index	168
Figure 76: Configuring Static VLAN Members by Interface	168
Figure 77: Configuring Static VLAN Members by Interface Range	169
Figure 78: Configuring Global Status of GVRP	171
Figure 79: Configuring GVRP for an Interface	171
Figure 80: Showing Dynamic VLANs Registered on the Switch	172
Figure 81: Showing the Members of a Dynamic VLAN	172
Figure 82: QinQ Operational Concept	173
Figure 83: Enabling QinQ Tunneling	177
Figure 84: Configuring CVLAN to SPVLAN Mapping Entries	179
Figure 85: Showing CVLAN to SPVLAN Mapping Entries	179
Figure 86: Adding an Interface to a QinQ Tunnel	180
Figure 87: Configuring Protocol VLANs	182
Figure 88: Displaying Protocol VLANs	183
Figure 89: Assigning Interfaces to Protocol VLANs	184
Figure 90: Showing the Interface to Protocol Group Mapping	185
Figure 91: Configuring IP Subnet VLANs	186
Figure 92: Showing IP Subnet VLANs	187
Figure 93: Configuring MAC-Based VLANs	188
Figure 94: Showing MAC-Based VLANs	189
Figure 95: Configuring VLAN Mirroring	190
Figure 96: Showing the VLANs to Mirror	190
Figure 97: Configuring VLAN Translation	191
Figure 98: Configuring VLAN Translation	192
Figure 99: Showing the Entries for VLAN Translation	192

Figure 100: Configuring MAC Address Learning	194
Figure 101: Configuring Static MAC Addresses	196
Figure 102: Displaying Static MAC Addresses	196
Figure 103: Setting the Address Aging Time	197
Figure 104: Displaying the Dynamic MAC Address Table	198
Figure 105: Clearing Entries in the Dynamic MAC Address Table	199
Figure 106: Mirroring Packets Based on the Source MAC Address	201
Figure 107: Showing the Source MAC Addresses to Mirror	201
Figure 108: Issuing MAC Address Traps (Global Configuration)	202
Figure 109: Issuing MAC Address Traps (Interface Configuration)	202
Figure 110: STP Root Ports and Designated Ports	204
Figure 111: MSTP Region, Internal Spanning Tree, Multiple Spanning Tree	204
Figure 112: Spanning Tree – Common Internal, Common, Internal	205
Figure 113: Configuring Port Loopback Detection	207
Figure 114: Configuring Global Settings for STA (STP)	211
Figure 115: Configuring Global Settings for STA (RSTP)	211
Figure 116: Configuring Global Settings for STA (MSTP)	212
Figure 117: Displaying Global Settings for STA	213
Figure 118: Determining the Root Port	215
Figure 119: Configuring Interface Settings for STA	217
Figure 120: STA Port Roles	219
Figure 121: Displaying Interface Settings for STA	220
Figure 122: Creating an MST Instance	222
Figure 123: Displaying MST Instances	222
Figure 124: Modifying the Priority for an MST Instance	223
Figure 125: Displaying Global Settings for an MST Instance	223
Figure 126: Adding a VLAN to an MST Instance	224
Figure 127: Displaying Members of an MST Instance	224
Figure 128: Configuring MSTP Interface Settings	226
Figure 129: Displaying MSTP Interface Settings	226
Figure 130: Configuring Rate Limits	228
Figure 131: Configuring Storm Control	230
Figure 132: Storm Control by Limiting the Traffic Rate	230
Figure 133: Storm Control by Shutting Down a Port	231
Figure 134: Configuring ATC Timers	232

Figure 135: Configuring ATC Interface Attributes	235
Figure 136: Setting the Default Port Priority	238
Figure 137: Setting the Queue Mode (Strict)	240
Figure 138: Setting the Queue Mode (WRR)	240
Figure 139: Setting the Queue Mode (Strict and WRR)	241
Figure 140: Mapping CoS Values to Egress Queues	243
Figure 141: Showing CoS Values to Egress Queue Mapping	243
Figure 142: Setting the Trust Mode	245
Figure 143: Configuring DSCP to DSCP Internal Mapping	247
Figure 144: Showing DSCP to DSCP Internal Mapping	247
Figure 145: Configuring CoS to DSCP Internal Mapping	249
Figure 146: Showing CoS to DSCP Internal Mapping	250
Figure 147: Configuring a Class Map	254
Figure 148: Showing Class Maps	254
Figure 149: Adding Rules to a Class Map	255
Figure 150: Showing the Rules for a Class Map	255
Figure 151: Configuring a Policy Map	263
Figure 152: Showing Policy Maps	263
Figure 153: Adding Rules to a Policy Map	264
Figure 154: Showing the Rules for a Policy Map	265
Figure 155: Attaching a Policy Map to a Port	266
Figure 156: Configuring a Voice VLAN	269
Figure 157: Configuring an OUI Telephony List	270
Figure 158: Showing an OUI Telephony List	270
Figure 159: Configuring Port Settings for a Voice VLAN	272
Figure 160: Configuring the Authentication Sequence	276
Figure 161: Authentication Server Operation	276
Figure 162: Configuring Remote Authentication Server (RADIUS)	279
Figure 163: Configuring Remote Authentication Server (TACACS+)	280
Figure 164: Configuring AAA Server Groups	280
Figure 165: Showing AAA Server Groups	281
Figure 166: Configuring Global Settings for AAA Accounting	283
Figure 167: Configuring AAA Accounting Methods	284
Figure 168: Showing AAA Accounting Methods	285
Figure 169: Configuring AAA Accounting Service for 802.1X Service	285

Figure 170: Configuring AAA Accounting Service for Command Service	286
Figure 171: Configuring AAA Accounting Service for Exec Service	286
Figure 172: Displaying a Summary of Applied AAA Accounting Methods	287
Figure 173: Displaying Statistics for AAA Accounting Sessions	287
Figure 174: Configuring AAA Authorization Methods	289
Figure 175: Showing AAA Authorization Methods	289
Figure 176: Configuring AAA Authorization Methods for Exec Service	290
Figure 177: Displaying the Applied AAA Authorization Method	290
Figure 178: Configuring User Accounts	292
Figure 179: Showing User Accounts	292
Figure 180: Configuring Global Settings for Web Authentication	294
Figure 181: Configuring Interface Settings for Web Authentication	295
Figure 182: Configuring Global Settings for Network Access	298
Figure 183: Configuring Interface Settings for Network Access	300
Figure 184: Configuring Link Detection for Network Access	301
Figure 185: Configuring a MAC Address Filter for Network Access	302
Figure 186: Showing the MAC Address Filter Table for Network Access	303
Figure 187: Showing Addresses Authenticated for Network Access	304
Figure 188: Configuring HTTPS	306
Figure 189: Downloading the Secure-Site Certificate	308
Figure 190: Configuring the SSH Server	311
Figure 191: Generating the SSH Host Key Pair	312
Figure 192: Showing the SSH Host Key Pair	313
Figure 193: Copying the SSH User's Public Key	314
Figure 194: Showing the SSH User's Public Key	315
Figure 195: Setting the Name of a Time Range	317
Figure 196: Showing a List of Time Ranges	318
Figure 197: Add a Rule to a Time Range	318
Figure 198: Showing the Rules Configured for a Time Range	319
Figure 199: Showing TCAM Utilization	320
Figure 200: Creating an ACL	321
Figure 201: Showing a List of ACLs	321
Figure 202: Configuring a Standard IPv4 ACL	323
Figure 203: Configuring an Extended IPv4 ACL	325
Figure 204: Configuring a Standard IPv6 ACL	326

Figure 205: Configuring an Extended IPv6 ACL	328
Figure 206: Configuring a MAC ACL	330
Figure 207: Configuring a ARP ACL	332
Figure 208: Binding a Port to an ACL	333
Figure 209: Configuring ACL Mirroring	334
Figure 210: Showing the VLANs to Mirror	335
Figure 211: Showing ACL Statistics	336
Figure 212: Configuring Global Settings for ARP Inspection	339
Figure 213: Configuring VLAN Settings for ARP Inspection	340
Figure 214: Configuring Interface Settings for ARP Inspection	341
Figure 215: Displaying Statistics for ARP Inspection	343
Figure 216: Displaying the ARP Inspection Log	344
Figure 217: Creating an IP Address Filter for Management Access	345
Figure 218: Showing IP Addresses Authorized for Management Access	346
Figure 219: Configuring Port Security	348
Figure 220: Configuring Port Authentication	349
Figure 221: Configuring Global Settings for 802.1X Port Authentication	351
Figure 222: Configuring Interface Settings for 802.1X Port Authenticator	355
Figure 223: Configuring Interface Settings for 802.1X Port Supplicant	357
Figure 224: Showing Statistics for 802.1X Port Authenticator	359
Figure 225: Showing Statistics for 802.1X Port Supplicant	360
Figure 226: Protecting Against DoS Attacks	362
Figure 227: Setting the Filter Type for IPv4 Source Guard	365
Figure 228: Configuring Static Bindings for IPv4 Source Guard	367
Figure 229: Displaying Static Bindings for IPv4 Source Guard	367
Figure 230: Showing the IPv4 Source Guard Binding Table	369
Figure 231: Setting the Filter Type for IPv6 Source Guard	371
Figure 232: Configuring Static Bindings for IPv6 Source Guard	373
Figure 233: Displaying Static Bindings for IPv6 Source Guard	373
Figure 234: Showing the IPv6 Source Guard Binding Table	375
Figure 235: Configuring Global Settings for DHCP Snooping	379
Figure 236: Configuring DHCP Snooping on a VLAN	380
Figure 237: Configuring the Port Mode for DHCP Snooping	381
Figure 238: Displaying the Binding Table for DHCP Snooping	382
Figure 239: Configuring Settings for System Memory Logs	385

Figure 240: Showing Error Messages Logged to System Memory	386
Figure 241: Configuring Settings for Remote Logging of Error Messages	387
Figure 242: Configuring SMTP Alert Messages	388
Figure 243: Configuring LLDP Timing Attributes	391
Figure 244: Configuring LLDP Interface Attributes	394
Figure 245: Configuring the Civic Address for an LLDP Interface	396
Figure 246: Showing the Civic Address for an LLDP Interface	396
Figure 247: Displaying Local Device Information for LLDP (General)	399
Figure 248: Displaying Local Device Information for LLDP (Port)	400
Figure 249: Displaying Local Device Information for LLDP (Port Details)	400
Figure 250: Displaying Remote Device Information for LLDP (Port)	407
Figure 251: Displaying Remote Device Information for LLDP (Port Details)	408
Figure 252: Displaying Remote Device Information for LLDP (End Node)	409
Figure 253: Displaying LLDP Device Statistics (General)	411
Figure 254: Displaying LLDP Device Statistics (Port)	411
Figure 255: Configuring Global Settings for SNMP	414
Figure 256: Configuring the Local Engine ID for SNMP	415
Figure 257: Configuring a Remote Engine ID for SNMP	416
Figure 258: Showing Remote Engine IDs for SNMP	416
Figure 259: Creating an SNMP View	418
Figure 260: Showing SNMP Views	418
Figure 261: Adding an OID Subtree to an SNMP View	419
Figure 262: Showing the OID Subtree Configured for SNMP Views	419
Figure 263: Creating an SNMP Group	424
Figure 264: Showing SNMP Groups	424
Figure 265: Setting Community Access Strings	425
Figure 266: Showing Community Access Strings	426
Figure 267: Configuring Local SNMPv3 Users	428
Figure 268: Showing Local SNMPv3 Users	428
Figure 269: Changing a Local SNMPv3 User Group	429
Figure 270: Configuring Remote SNMPv3 Users	431
Figure 271: Showing Remote SNMPv3 Users	431
Figure 272: Configuring Trap Managers (SNMPv1)	435
Figure 273: Configuring Trap Managers (SNMPv2c)	435
Figure 274: Configuring Trap Managers (SNMPv3)	435

Figure 275: Showing Trap Managers	436
Figure 276: Creating SNMP Notification Logs	437
Figure 277: Showing SNMP Notification Logs	438
Figure 278: Showing SNMP Statistics	439
Figure 279: Configuring an RMON Alarm	442
Figure 280: Showing Configured RMON Alarms	442
Figure 281: Configuring an RMON Event	444
Figure 282: Showing Configured RMON Events	445
Figure 283: Configuring an RMON History Sample	446
Figure 284: Showing Configured RMON History Samples	447
Figure 285: Showing Collected RMON History Samples	447
Figure 286: Configuring an RMON Statistical Sample	449
Figure 287: Showing Configured RMON Statistical Samples	449
Figure 288: Showing Collected RMON Statistical Samples	450
Figure 289: Configuring a Switch Cluster	452
Figure 290: Configuring a Cluster Members	453
Figure 291: Showing Cluster Members	453
Figure 292: Showing Cluster Candidates	454
Figure 293: Managing a Cluster Member	455
Figure 294: ERPS Ring Components	456
Figure 295: Ring Interconnection Architecture (Multi-ring/Ladder Network)	458
Figure 296: Setting ERPS Global Status	460
Figure 297: Sub-ring with Virtual Channel	469
Figure 298: Sub-ring without Virtual Channel	470
Figure 299: Non-ERPS Device Protection	471
Figure 300: Creating an ERPS Ring	474
Figure 301: Creating an ERPS Ring	475
Figure 302: Showing Configured ERPS Rings	475
Figure 303: Blocking an ERPS Ring Port	480
Figure 304: Single CFM Maintenance Domain	481
Figure 305: Multiple CFM Maintenance Domains	482
Figure 306: Configuring Global Settings for CFM	486
Figure 307: Configuring Interfaces for CFM	487
Figure 308: Configuring Maintenance Domains	491
Figure 309: Showing Maintenance Domains	491

Figure 310: Configuring Detailed Settings for Maintenance Domains	492
Figure 311: Creating Maintenance Associations	495
Figure 312: Showing Maintenance Associations	495
Figure 313: Configuring Detailed Settings for Maintenance Associations	496
Figure 314: Configuring Maintenance End Points	497
Figure 315: Showing Maintenance End Points	498
Figure 316: Configuring Remote Maintenance End Points	499
Figure 317: Showing Remote Maintenance End Points	500
Figure 318: Transmitting Link Trace Messages	502
Figure 319: Transmitting Loopback Messages	503
Figure 320: Transmitting Delay-Measure Messages	505
Figure 321: Showing Information on Local MEPs	506
Figure 322: Showing Detailed Information on Local MEPs	508
Figure 323: Showing Information on Local MIPs	509
Figure 324: Showing Information on Remote MEPs	510
Figure 325: Showing Detailed Information on Remote MEPs	512
Figure 326: Showing the Link Trace Cache	514
Figure 327: Showing Settings for the Fault Notification Generator	515
Figure 328: Showing Continuity Check Errors	516
Figure 329: Enabling OAM for Local Ports	519
Figure 330: Displaying Statistics for OAM Messages	520
Figure 331: Displaying the OAM Event Log	521
Figure 332: Displaying Status of Remote Interfaces	522
Figure 333: Running a Remote Loop Back Test	524
Figure 334: Displaying the Results of Remote Loop Back Testing	525
Figure 335: Configuring UDLD Protocol Intervals	527
Figure 336: Configuring UDLD Interface Settings	529
Figure 337: Displaying UDLD Neighbor Information	530
Figure 338: Multicast Filtering Concept	532
Figure 339: IGMP Protocol	533
Figure 340: Configuring General Settings for IGMP Snooping	539
Figure 341: Configuring a Static Interface for a Multicast Router	540
Figure 342: Showing Static Interfaces Attached a Multicast Router	541
Figure 343: Showing Current Interfaces Attached a Multicast Router	541
Figure 344: Assigning an Interface to a Multicast Service	542

Figure 345: Showing Static Interfaces Assigned to a Multicast Service	543
Figure 346: Showing Current Interfaces Attached a Multicast Router	543
Figure 347: Configuring IGMP Snooping on a VLAN	549
Figure 348: Showing Interface Settings for IGMP Snooping	549
Figure 349: Dropping IGMP Query or Multicast Data Packets	550
Figure 350: Showing Multicast Groups Learned by IGMP Snooping	551
Figure 351: Displaying IGMP Snooping Statistics – Query	554
Figure 352: Displaying IGMP Snooping Statistics – VLAN	555
Figure 353: Displaying IGMP Snooping Statistics – Port	555
Figure 354: Enabling IGMP Filtering and Throttling	557
Figure 355: Creating an IGMP Filtering Profile	558
Figure 356: Showing the IGMP Filtering Profiles Created	558
Figure 357: Adding Multicast Groups to an IGMP Filtering Profile	559
Figure 358: Showing the Groups Assigned to an IGMP Filtering Profile	559
Figure 359: Configuring IGMP Filtering and Throttling Interface Settings	561
Figure 360: Configuring General Settings for MLD Snooping	563
Figure 361: Configuring Immediate Leave for MLD Snooping	564
Figure 362: Configuring a Static Interface for an IPv6 Multicast Router	565
Figure 363: Showing Static Interfaces Attached an IPv6 Multicast Router	565
Figure 364: Showing Current Interfaces Attached an IPv6 Multicast Router	565
Figure 365: Assigning an Interface to an IPv6 Multicast Service	567
Figure 366: Showing Static Interfaces Assigned to an IPv6 Multicast Service	567
Figure 367: Showing Current Interfaces Assigned to an IPv6 Multicast Service	568
Figure 368: Showing IPv6 Multicast Services and Corresponding Sources	569
Figure 369: IGMP Proxy Routing	570
Figure 370: Configuring IGMP Proxy Routing	572
Figure 371: Configuring IGMP Interface Settings	575
Figure 372: Configuring Static IGMP Groups	576
Figure 373: Showing Static IGMP Groups	576
Figure 374: Displaying Multicast Groups Learned from IGMP (Information)	579
Figure 375: Displaying Multicast Groups Learned from IGMP (Details)	579
Figure 376: MVR Concept	580
Figure 377: Configuring Global Settings for MVR	582
Figure 378: Configuring Domain Settings for MVR	584
Figure 379: Configuring an MVR Group Address Profile	585

Figure 380: Displaying MVR Group Address Profiles	585
Figure 381: Assigning an MVR Group Address Profile to a Domain	586
Figure 382: Showing the MVR Group Address Profiles Assigned to a Domain	586
Figure 383: Configuring Interface Settings for MVR	589
Figure 384: Assigning Static MVR Groups to an Interface	590
Figure 385: Showing the Static MVR Groups Assigned to a Port	591
Figure 386: Displaying MVR Receiver Groups	592
Figure 387: Displaying MVR Statistics – Query	594
Figure 388: Displaying MVR Statistics – VLAN	595
Figure 389: Displaying MVR Statistics – Port	596
Figure 390: Configuring Global Settings for MVR6	599
Figure 391: Configuring Domain Settings for MVR6	600
Figure 392: Configuring an MVR6 Group Address Profile	602
Figure 393: Displaying MVR6 Group Address Profiles	602
Figure 394: Assigning an MVR6 Group Address Profile to a Domain	603
Figure 395: Showing MVR6 Group Address Profiles Assigned to a Domain	603
Figure 396: Configuring Interface Settings for MVR6	605
Figure 397: Assigning Static MVR6 Groups to a Port	606
Figure 398: Showing the Static MVR6 Groups Assigned to a Port	607
Figure 399: Displaying MVR6 Receiver Groups	608
Figure 400: Displaying MVR6 Statistics – Query	610
Figure 401: Displaying MVR6 Statistics – VLAN	611
Figure 402: Displaying MVR6 Statistics – Port	612
Figure 403: Configuring a Static IPv4 Address	615
Figure 404: Configuring a Dynamic IPv4 Address	616
Figure 405: Showing the Configured IPv4 Address for an Interface	617
Figure 406: Send DHCP Inform Requests for Additional Information	618
Figure 407: Configuring the IPv6 Default Gateway	619
Figure 408: Configuring General Settings for an IPv6 Interface	623
Figure 409: Configuring RA Guard for an IPv6 Interface	623
Figure 410: Configuring an IPv6 Address	626
Figure 411: Showing Configured IPv6 Addresses	627
Figure 412: Showing IPv6 Neighbors	629
Figure 413: Showing IPv6 Statistics (IPv6)	633
Figure 414: Showing IPv6 Statistics (ICMPv6)	634

Figure 415: Showing IPv6 Statistics (UDP)	634
Figure 416: Showing Reported MTU Values	635
Figure 417: Configuring General Settings for DNS	638
Figure 418: Configuring a List of Domain Names for DNS	639
Figure 419: Showing the List of Domain Names for DNS	639
Figure 420: Configuring a List of Name Servers for DNS	640
Figure 421: Showing the List of Name Servers for DNS	641
Figure 422: Configuring Static Entries in the DNS Table	641
Figure 423: Showing Static Entries in the DNS Table	642
Figure 424: Showing Entries in the DNS Cache	643
Figure 425: Specifying a DHCP Client Identifier	645
Figure 426: Layer 3 DHCP Relay Service	645
Figure 427: Configuring DHCP Relay Service	646
Figure 428: DHCP Server	646
Figure 429: Enabling the DHCP Server	647
Figure 430: Configuring Excluded Addresses on the DHCP Server	648
Figure 431: Showing Excluded Addresses on the DHCP Server	649
Figure 432: Configuring DHCP Server Address Pools (Network)	651
Figure 433: Configuring DHCP Server Address Pools (Host)	652
Figure 434: Showing Configured DHCP Server Address Pools	652
Figure 435: Shows Addresses Assigned by the DHCP Server	653
Figure 436: Enabling the UDP Helper	654
Figure 437: Specifying UDP Destination Ports	655
Figure 438: Showing the UDP Destination Ports	655
Figure 439: Specifying the Target Server or Subnet for UDP Requests	657
Figure 440: Showing the Target Server or Subnet for UDP Requests	657
Figure 441: Configuring Global Settings for PPPoE Intermediate Agent	658
Figure 442: Configuring Interface Settings for PPPoE Intermediate Agent	660
Figure 443: Showing PPPoE Intermediate Agent Statistics	661
Figure 444: Virtual Interfaces and Layer 3 Routing	664
Figure 445: Pinging a Network Device	668
Figure 446: Tracing the Route to a Network Device	669
Figure 447: Proxy ARP	671
Figure 448: Configuring General Settings for ARP	672
Figure 449: Configuring Static ARP Entries	673

Figure 450: Displaying Static ARP Entries	673
Figure 451: Displaying ARP Entries	674
Figure 452: Displaying ARP Statistics	675
Figure 453: Configuring Static Routes	676
Figure 454: Displaying Static Routes	677
Figure 455: Displaying the Routing Table	678
Figure 456: Setting the Maximum ECMP Number	680
Figure 457: Master Virtual Router with Backup Routers	681
Figure 458: Several Virtual Master Routers Using Backup Routers	682
Figure 459: Several Virtual Master Routers Configured for Mutual Backup and Load Sharing	682
Figure 460: Configuring the VRRP Group ID	686
Figure 461: Showing Configured VRRP Groups	686
Figure 462: Setting the Virtual Router Address for a VRRP Group	687
Figure 463: Showing the Virtual Addresses Assigned to VRRP Groups	687
Figure 464: Configuring Detailed Settings for a VRRP Group	688
Figure 465: Showing Counters for Errors Found in VRRP Packets	689
Figure 466: Showing Counters for Errors Found in a VRRP Group	690
Figure 467: Configuring RIP	692
Figure 468: Configuring General Settings for RIP	696
Figure 469: Clearing Entries from the Routing Table	697
Figure 470: Adding Network Interfaces to RIP	698
Figure 471: Showing Network Interfaces Using RIP	699
Figure 472: Specifying a Passive RIP Interface	700
Figure 473: Showing Passive RIP Interfaces	700
Figure 474: Specifying a Static RIP Neighbor	701
Figure 475: Showing Static RIP Neighbors	701
Figure 476: Redistributing External Routes into RIP	702
Figure 477: Showing External Routes Redistributed into RIP	703
Figure 478: Setting the Distance Assigned to External Routes	704
Figure 479: Showing the Distance Assigned to External Routes	704
Figure 480: Configuring a Network Interface for RIP	707
Figure 481: Showing RIP Network Interface Settings	708
Figure 482: Showing RIP Interface Settings	708
Figure 483: Showing RIP Peer Information	709
Figure 484: Resetting RIP Statistics	710

Figure 485: Configuring OSPF	710
Figure 486: OSPF Areas	712
Figure 487: Defining OSPF Network Areas Based on Addresses	713
Figure 488: Showing OSPF Network Areas	714
Figure 489: Showing OSPF Process Identifiers	714
Figure 490: AS Boundary Router	716
Figure 491: Configure General Settings for OSPF	717
Figure 492: Showing General Settings for OSPF	718
Figure 493: Adding an NSSA or Stub	719
Figure 494: Showing NSSAs or Stubs	720
Figure 495: OSPF NSSA	720
Figure 496: Configuring Protocol Settings for an NSSA	723
Figure 497: OSPF Stub Area	723
Figure 498: Configuring Protocol Settings for a Stub	725
Figure 499: Displaying Information on NSSA and Stub Areas	726
Figure 500: Route Summarization for ABRs	726
Figure 501: Configuring Route Summaries for an Area Range	727
Figure 502: Showing Configured Route Summaries	728
Figure 503: Redistributing External Routes	728
Figure 504: Importing External Routes	730
Figure 505: Showing Imported External Route Types	730
Figure 506: Summarizing External Routes	731
Figure 507: Showing Summary Addresses for External Routes	732
Figure 508: Configuring Settings for All Interfaces Assigned to a VLAN	735
Figure 509: Configuring Settings for a Specific Area Assigned to a VLAN	736
Figure 510: Showing OSPF Interfaces	737
Figure 511: Showing MD5 Authentication Keys	737
Figure 512: OSPF Virtual Link	738
Figure 513: Adding a Virtual Link	739
Figure 514: Showing Virtual Links	739
Figure 515: Configuring Detailed Settings for a Virtual Link	740
Figure 516: Showing MD5 Authentication Keys	740
Figure 517: Displaying Information in the Link State Database	742
Figure 518: Displaying Neighbor Routers Stored in the Link State Database	744
Figure 519: Enabling IPv4 Multicast Routing	748

Figure 520: Enabling IPv6 Multicast Routing	749
Figure 521: Displaying the IPv4 Multicast Routing Table	751
Figure 522: Displaying Detailed Entries from IPv4 Multicast Routing Table	752
Figure 523: Displaying the IPv6 Multicast Routing Table	752
Figure 524: Displaying Detailed Entries from IPv6 Multicast Routing Table	753
Figure 525: Enabling PIM Multicast Routing	754
Figure 526: Configuring PIM Interface Settings (Dense Mode)	758
Figure 527: Configuring PIM Interface Settings (Sparse Mode)	759
Figure 528: Showing PIM Neighbors	759
Figure 529: Configuring Global Settings for PIM-SM	761
Figure 530: Configuring a PIM-SM BSR Candidate	762
Figure 531: Configuring a PIM Static Rendezvous Point	764
Figure 532: Showing PIM Static Rendezvous Points	764
Figure 533: Configuring a PIM RP Candidate	766
Figure 534: Showing Settings for a PIM RP Candidate	766
Figure 535: Showing Information About the PIM BSR	767
Figure 536: Showing PIM RP Mapping	768
Figure 537: Enabling PIMv6 Multicast Routing	769
Figure 538: Configuring PIMv6 Interface Settings (Dense Mode)	774
Figure 539: Configuring PIMv6 Interface Settings (Sparse Mode)	774
Figure 540: Showing PIMv6 Neighbors	775
Figure 541: Configuring Global Settings for PIM6-SM	777
Figure 542: Configuring a PIM6-SM BSR Candidate	778
Figure 543: Configuring a PIM6 Static Rendezvous Point	779
Figure 544: Showing PIM6 Static Rendezvous Points	780
Figure 545: Configuring a PIM6 RP Candidate	781
Figure 546: Showing Settings for a PIM6 RP Candidate	782
Figure 547: Showing Information About the PIM6 BSR	783
Figure 548: Showing PIM6 RP Mapping	784

Tables

Table 1: Key Features	39
Table 2: System Defaults	47
Table 3: Web Page Configuration Buttons	55
Table 4: Switch Main Menu	56
Table 5: Predefined Summer-Time Parameters	102
Table 6: Port Statistics	127
Table 7: LACP Port Counters	145
Table 8: LACP Internal Configuration Information	146
Table 9: LACP Remote Device Configuration Information	148
Table 10: Traffic Segmentation Forwarding	154
Table 11: Recommended STA Path Cost Range	214
Table 12: Default STA Path Costs	215
Table 13: IEEE 802.1p Egress Queue Priority Mapping	241
Table 14: CoS Priority Levels	242
Table 15: Mapping Internal Per-hop Behavior to Hardware Queues	242
Table 16: Default Mapping of DSCP Values to Internal PHB/Drop Values	246
Table 17: Default Mapping of CoS/CFI to Internal PHB/Drop Precedence	248
Table 18: Dynamic QoS Profiles	296
Table 19: HTTPS System Support	305
Table 20: ARP Inspection Statistics	342
Table 21: ARP Inspection Log	343
Table 22: 802.1X Statistics	358
Table 23: Logging Levels	384
Table 24: LLDP MED Location CA Types	395
Table 25: Chassis ID Subtype	397
Table 26: System Capabilities	397
Table 27: Port ID Subtype	398
Table 28: Remote Port Auto-Negotiation Advertised Capability	402
Table 29: SNMPv3 Security Models and Levels	412

Table 30: Supported Notification Messages	421
Table 31: ERPS Request/State Priority	477
Table 32: Remote MEP Priority Levels	489
Table 33: MEP Defect Descriptions	489
Table 34: OAM Operation State	517
Table 35: Remote Loopback Status	523
Table 36: Show IPv6 Neighbors - display description	628
Table 37: Show IPv6 Statistics - display description	630
Table 38: Show MTU - display description	635
Table 39: Options 60, 66 and 67 Statements	644
Table 40: Options 55 and 124 Statements	644
Table 41: Address Resolution Protocol	670
Table 42: ARP Statistics	674
Table 43: VRRP Group Statistics	689
Table 44: OSPF System Information	717
Table 45: Troubleshooting Chart	793

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 39](#)

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 1: Key Features

Feature	Description
Configuration Backup and Restore	Using management station or FTP/TFTP server
Authentication	Console, Telnet, web – user name/password, RADIUS, TACACS+ Port – IEEE 802.1X, MAC address filtering SNMP v1/2c - Community strings SNMP version 3 – MD5 or SHA password Telnet – SSH Web – HTTPS
General Security Measures	AAA ARP Inspection DHCP Snooping (with Option 82 relay information) DoS Protection IP Source Guard PPPoE Intermediate Agent Port Authentication – IEEE 802.1X Port Security – MAC address filtering
Access Control Lists	Supports up to 512 ACLs, 2048 rules per ACL, and 2048 rules per system
DHCP/DHCPv6	Client, Relay, Relay Option 82, Server
DNS	Client and Proxy service
Port Configuration	Speed, duplex mode, and flow control
Port Trunking	Supports up to 16 trunks per switch (32 per stack) – static or dynamic trunking (LACP)
Port Mirroring	27 sessions, across switch or stack, one or more source ports to one analysis port
Congestion Control	Rate Limiting Throttling for broadcast, multicast, unknown unicast storms

Table 1: Key Features (Continued)

Feature	Description
Address Table	16K MAC addresses in the forwarding table, 1K static MAC addresses; 1760 entries in the ARP cache, 256 static ARP entries, 3836 dynamic ARP entries; 512 static IP routes, 512 IP interfaces; 3996 IPv4 entries in the host table; 12900 IPv4 entries in routing table; 1K L2 IPv4 multicast groups (shared with MAC table); 1K L3 IPv4 multicast groups; 2997 IPv6 entries in the host table; 2996 IPv6 entries in routing table; 1K L3 IPv6 multicast groups
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 4093 using IEEE 802.1Q, port-based, protocol-based, voice VLANs, and QinQ tunnel
Traffic Prioritization	Default port priority, traffic class map, queue scheduling, IP Precedence, or Differentiated Services Code Point (DSCP)
Quality of Service	Supports Differentiated Services (DiffServ)
Link Layer Discovery Protocol	Used to discover basic information about neighboring devices
Switch Clustering	Supports up to 36 member switches in a cluster
Connectivity Fault Management	Connectivity monitoring using continuity check messages, fault verification through loop back messages, and fault isolation by examining end-to-end connections (IEEE 802.1ag)
ERPS	Supports Ethernet Ring Protection Switching for increased availability of Ethernet rings (G.8032)
IP Routing	Routing Information Protocol (RIPv2), 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, MLD snooping and query for Layer 3, IGMP for Layer 3, and Multicast VLAN Registration
Multicast Routing	Protocol-Independent Multicasting - Dense Mode and Sparse Mode (PIM-DM, PIM-SM, PIM6-DM, PIM6-SM)
Remote Device Management	Supports Ethernet OAM functions for attached CPEs (IEEE 802.3ah, ITU-T Y.1731)

* 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. Storm suppression prevents broadcast, multicast, and unknown unicast 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+). Port-based authentication is also supported via the IEEE 802.1X protocol. This protocol uses Extensible Authentication Protocol over LANs (EAPOL) to request user credentials from the 802.1X client, and then uses the EAP between the switch and the authentication server to verify the client's right to access the network via an authentication server (i.e., RADIUS or TACACS+ server).

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. While DHCP snooping is provided to prevent malicious attacks from insecure ports. While PPPoE Intermediate Agent supports authentication of a client for a service provider.

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 A DHCP server is provided to assign IP addresses to host devices. Since DHCP uses a broadcast mechanism, a DHCP server and its client must physically reside on the same subnet. Since it is not practical to have a DHCP server on every subnet, DHCP Relay is also supported to allow dynamic configuration of local clients from a DHCP

server located in a different network. And 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).

Rate Limiting This feature controls the maximum rate for traffic transmitted or received 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.

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 16 trunks per switch and 32 per stack.

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 16K 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 12 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).

Connectivity Fault Management The switch provides connectivity fault monitoring for end-to-end connections within a designated service area by using continuity check messages which can detect faults in maintenance points, fault verification through loop back messages, and fault isolation with link trace messages.

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 dynamically learned via GVRP, or ports 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.
- ◆ Use private VLANs to restrict traffic to pass only between data ports and the uplink ports, thereby isolating adjacent ports within the same VLAN, and allowing you to limit the total number of VLANs that need to be configured.
- ◆ Use protocol VLANs to restrict traffic to specified interfaces based on protocol type.

IEEE 802.1Q Tunneling (QinQ) This feature 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.

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) scheduling, 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. 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.

Ethernet Ring Protection Switching ERPS can be used to increase the availability and robustness of Ethernet rings, such as those used in Metropolitan Area Networks (MAN). ERPS provides Layer 2 loop avoidance and fast reconvergence in Layer 2 ring topologies, supporting up to 255 nodes in the ring structure. It can also function with IEEE 802.1ag to support link monitoring when non-participating devices exist within the Ethernet ring.

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, and Routing Information Protocol (RIP).

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

RIP – This protocol uses a distance-vector 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.

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.

Operation, Administration, and Maintenance

The switch provides OAM remote management tools required to monitor and maintain the links to subscriber CPEs (Customer Premise Equipment). This section describes functions including enabling OAM for selected ports, loopback testing, and displaying remote device information.

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, MLD Snooping and Query for IPv6, and IGMP at Layer 3 to manage multicast group registration. It also supports Multicast VLAN Registration

(MVR for IPv4 and MVR6 for IPv6) which allows common multicast traffic, such as television channels, to be transmitted across a single network-wide multicast VLAN shared by hosts residing in other standard or private VLAN groups, while preserving security and data isolation for normal traffic.

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 2: System Defaults

Function	Parameter	Default
Console Port Connection	Baud Rate	115200 bps
	Data bits	8
	Stop bits	1
	Parity	none
	Local Console Timeout	600 seconds

Table 2: System Defaults (Continued)

Function	Parameter	Default
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
	802.1X Port Authentication	Disabled
	Web Authentication	Disabled
	MAC Authentication	Disabled
	PPPoE Intermediate Agent	Disabled
	HTTPS	Enabled
	SSH	Disabled
	Port Security	Disabled
	IP Filtering	Disabled
	DHCP Snooping	Disabled
	IP Source Guard	Disabled (all ports)
Web Management	HTTP Server	Enabled
	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)
Port Configuration	Admin Status	Enabled
	Auto-negotiation	Enabled
	Flow Control	Disabled
Port Trunking	Static Trunks	None
	LACP (all ports)	Disabled

Table 2: System Defaults (Continued)

Function	Parameter	Default
Congestion Control	Rate Limiting	Disabled
	Storm Control	Broadcast: Enabled (64 kbits/sec) Multicast: Disabled Unknown Unicast: Disabled
	Auto Traffic Control	Disabled
	Address Table	Aging Time 300 seconds
Spanning Tree Algorithm	Status	Enabled, RSTP (Defaults: RSTP standard)
	Edge Ports	Disabled
LLDP	Status	Enabled
ERPS	Status	Disabled
CFM	Status	Enabled
OAM	Status	Disabled
Virtual LANs	Default VLAN	1
	PVID	1
	Acceptable Frame Type	All
	Ingress Filtering	Disabled
	Switchport Mode (Egress Mode)	Hybrid
	GVRP (global)	Disabled
	GVRP (port interface)	Disabled
	QinQ Tunneling	Disabled
	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 DSCP Priority Disabled
	IP Settings	Management. VLAN VLAN 1 IP Address DHCP assigned Subnet Mask 255.255.255.0 Default Gateway Not configured DHCP Client: Enabled DNS Proxy service: Disabled BOOTP Disabled

Table 2: System Defaults (Continued)

Function	Parameter	Default
Unicast Routing	ARP	Enabled Cache Timeout: 20 minutes Proxy: Disabled
	RIP	Disabled
	OSPF	Disabled
	OSPFv3	Disabled
	BGPv4	Disabled
Multicast Filtering	IGMP Snooping (Layer 2)	Snooping: Enabled Querier: Disabled
	MLD Snooping (Layer 2 IPv6)	Snooping: Enabled Querier: Disabled
	Multicast VLAN Registration	Disabled
	IGMP Proxy Reporting	Disabled
	IGMP (Layer 3) IGMP Proxy (Layer 3)	Disabled Disabled
System Log	Status	Enabled
	Messages Logged to RAM	Levels 0-7 (all)
	Messages Logged to Flash	Levels 0-3
SMTP Email Alerts	Event Handler	Enabled (but no server defined)
SNTP	Clock Synchronization	Disabled
Switch Clustering	Status	Disabled
	Commander	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 53](#)
- ◆ ["Basic Management Tasks" on page 79](#)
- ◆ ["Interface Configuration" on page 115](#)
- ◆ ["VLAN Configuration" on page 159](#)
- ◆ ["Address Table Settings" on page 193](#)
- ◆ ["Spanning Tree Algorithm" on page 203](#)
- ◆ ["Congestion Control" on page 227](#)
- ◆ ["Class of Service" on page 237](#)
- ◆ ["Quality of Service" on page 251](#)
- ◆ ["VoIP Traffic Configuration" on page 267](#)
- ◆ ["Security Measures" on page 273](#)
- ◆ ["Basic Administration Protocols" on page 383](#)
- ◆ ["Multicast Filtering" on page 531](#)
- ◆ ["IP Configuration" on page 613](#)
- ◆ ["IP Services" on page 637](#)
- ◆ ["General IP Routing" on page 663](#)

- ◆ ["Configuring Router Redundancy" on page 681](#)
- ◆ ["Unicast Routing" on page 691](#)
- ◆ ["Multicast Routing" on page 745](#)

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 6, Mozilla Firefox 4, or Google Chrome 29, or more recent 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, BOOTP or DHCP protocol. (See “Initial Switch Configuration” in 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 [“Configuring User Accounts” on page 290](#).)
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 213](#).

Note: Users are automatically logged off of the HTTP server or HTTPS server if no input is detected for 600 seconds.

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.” The administrator has full access privileges to configure any parameters in the web interface. The default user name and password for guest access is “guest.” The guest only has read access for most configuration parameters. Refer to [“Configuring User Accounts” on page 290](#) for more details.

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 GTL-2891 Gigabit Ethernet switch. 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 levelone 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 3: 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

GTL-2891



NOTE: If stacking is enabled, stacking ports 25/26 are not displayed on the panel graphic.

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 4: Switch Main Menu

Menu	Description	Page
System		
General	Provides basic system description, including contact information	80
Switch	Shows the number of ports, hardware version, power status, and firmware version numbers	81
Capability	Enables support for jumbo frames; shows the bridge extension parameters	82 , 83
File		85
Copy	Allows the transfer and copying files	85
Set Startup	Sets the startup file	88
Show	Shows the files stored in flash memory; allows deletion of files	89
Automatic Operation Code Upgrade	Automatically upgrades operation code if a newer version is found on the server	89
Time		93
Configure General		
Manual	Manually sets the current time	94
SNTP	Configures SNTP polling interval	95
NTP	Configures NTP authentication parameters	95
Configure Time Server	Configures a list of SNTP servers	96
Configure SNTP Server	Sets the IP address for SNTP time servers	96
Add NTP Server	Adds NTP time server and index of authentication key	97
Show NTP Server	Shows list of configured NTP time servers	97
Add NTP Authentication Key	Adds key index and corresponding MD5 key	99
Show NTP Authentication Key	Shows list of configured authentication keys	99
Configure Time Zone	Sets the local time zone for the system clock	100
Configure Summer Time	Configures summer time settings	101
Console	Sets console port connection parameters	103
Telnet	Sets Telnet connection parameters	105
CPU Utilization	Displays information on CPU utilization	106
Memory Status	Shows memory utilization parameters	107
Stacking	Configure stacking functions	108
Configure Master Button	Set Master unit for stack	108
Configure Stacking Button	Enable stacking on 10G ports	109

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Renumber	Reset stack identification numbers	110
Reset	Restarts the switch immediately, at a specified time, after a specified delay, or at a periodic interval	111
Interface		115
Port		116
General		116
Configure by Port List	Configures connection settings per port	116
Configure by Port Range	Configures connection settings for a range of ports	118
Show Information	Displays port connection status	119
Mirror		120
Add	Sets the source and target ports for mirroring	120
Show	Shows the configured mirror sessions	120
Statistics	Shows Interface, Etherlike, and RMON port statistics	126
Chart	Shows Interface, Etherlike, and RMON port statistics	126
Transceiver	Shows identifying information and operational parameters for optical transceivers which support Digital Diagnostic Monitoring (DDM), and configures thresholds for alarm and warning messages for optical transceivers which support DDM	131 132
Cable Test	Performs cable diagnostics for selected port to diagnose any cable faults (short, open etc.) and report the cable length	134
Trunk		135
Static		137
Configure Trunk		137
Add	Creates a trunk, along with the first port member	137
Show	Shows the configured trunk identifiers	137
Add Member	Specifies ports to group into static trunks	137
Show Member	Shows the port members for the selected trunk	137
Configure General		137
Configure	Configures trunk connection settings	137
Show Information	Displays trunk connection settings	137
Dynamic		139
Configure Aggregator	Configures administration key and timeout for specific LACP groups	139
Configure Aggregation Port		137
Configure		137
General	Allows ports to dynamically join trunks	139

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Actor	Configures parameters for link aggregation group members on the local side	139
Partner	Configures parameters for link aggregation group members on the remote side	139
Show Information		145
Counters	Displays statistics for LACP protocol messages	145
Internal	Displays configuration settings and operational state for the local side of a link aggregation	146
Neighbors	Displays configuration settings and operational state for the remote side of a link aggregation	148
Configure Trunk		139
Configure	Configures connection settings	139
Show	Displays port connection status	139
Show Member	Shows the active members in a trunk	139
Statistics	Shows Interface, Etherlike, and RMON port statistics	126
Chart	Shows Interface, Etherlike, and RMON port statistics	126
Load Balance	Sets the load-distribution method among ports in aggregated links	149
Green Ethernet	Adjusts the power provided to ports based on the length of the cable used to connect to other devices	151
RSPAN	Mirrors traffic from remote switches for analysis at a destination port on the local switch	122
Traffic Segmentation		153
Configure Global	Enables traffic segmentation globally	153
Configure Session	Configures the uplink and down-link ports for a segmented group of ports	154
VLAN Trunking	Allows unknown VLAN groups to pass through the specified interface	156
VLAN	Virtual LAN	159
Static		
Add	Creates VLAN groups	162
Show	Displays configured VLAN groups	162
Modify	Configures group name and administrative status	162
Edit Member by VLAN	Specifies VLAN attributes per VLAN	165
Edit Member by Interface	Specifies VLAN attributes per interface	165
Edit Member by Interface Range	Specifies VLAN attributes per interface range	165
Dynamic		
Configure General	Enables GVRP VLAN registration protocol globally	169
Configure Interface	Configures GVRP status and timers per interface	169

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show Dynamic VLAN		169
Show VLAN	Shows the VLANs this switch has joined through GVRP	169
Show VLAN Member	Shows the interfaces assigned to a VLAN through GVRP	169
Tunnel	IEEE 802.1Q (QinQ) Tunneling	172
Configure Global	Sets tunnel mode for the switch	176
Configure Service	Sets a CVLAN to SPVLAN mapping entry	178
Configure Interface	Sets the tunnel mode for any participating interface	179
Protocol		181
Configure Protocol		181
Add	Creates a protocol group, specifying supported protocols	181
Show	Shows configured protocol groups	181
Configure Interface		183
Add	Maps a protocol group to a VLAN	183
Show	Shows the protocol groups mapped to each VLAN	183
IP Subnet		185
Add	Maps IP subnet traffic to a VLAN	185
Show	Shows IP subnet to VLAN mapping	185
MAC-Based		187
Add	Maps traffic with specified source MAC address to a VLAN	187
Show	Shows source MAC address to VLAN mapping	187
Mirror		189
Add	Mirrors traffic from one or more source VLANs to a target port	189
Show	Shows mirror list	189
Translation		191
Add	Maps VLAN IDs between the customer and service provider	191
Show	Displays the configuration settings for VLAN translation	191
MAC Address		193
Learning Status	Enables MAC address learning on selected interfaces	193
Static		195
Add	Configures static entries in the address table	195
Show	Displays static entries in the address table	195
Dynamic		
Configure Aging	Sets timeout for dynamically learned entries	197
Show Dynamic MAC	Displays dynamic entries in the address table	197

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
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	199
Mirror	Mirrors traffic matching a specified source address from any port on the switch to a target port	200
MAC Notification		201
Configure Global	Issues a trap when a dynamic MAC address is added or removed.	201
Configure Interface	Enables MAC authentication traps on the current interface.	201
Spanning Tree		203
Loopback Detection	Configures Loopback Detection parameters	205
STA	Spanning Tree Algorithm	
Configure Global		
Configure	Configures global bridge settings for STP, RSTP and MSTP	207
Show Information	Displays STA values used for the bridge	212
Configure Interface		
Configure	Configures interface settings for STA	213
Show Information	Displays interface settings for STA	218
MSTP	Multiple Spanning Tree Algorithm	220
Configure Global		220
Add	Configures initial VLAN and priority for an MST instance	220
Modify	Configures the priority or an MST instance	220
Show	Configures global settings for an MST instance	220
Add Member	Adds VLAN members for an MST instance	220
Show Member	Adds or deletes VLAN members for an MST instance	220
Show Information	Displays MSTP values used for the bridge	
Configure Interface		225
Configure	Configures interface settings for an MST instance	225
Show Information	Displays interface settings for an MST instance	225
Traffic		
Rate Limit	Sets the input and output rate limits for a port	227
Storm Control	Sets the broadcast storm threshold for each interface	228
Auto Traffic Control	Sets thresholds for broadcast and multicast storms which can be used to trigger configured rate limits or to shut down a port	230
Configure Global	Sets the time to apply the control response after traffic has exceeded the upper threshold, and the time to release the control response after traffic has fallen beneath the lower threshold	231

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure Interface	Sets the storm control mode (broadcast or multicast), the traffic thresholds, the control response, to automatically release a response of rate limiting, or to send related SNMP trap messages	233
Priority		
Default Priority	Sets the default priority for each port or trunk	237
Queue	Sets queue mode for the switch; sets the service weight for each queue that will use a weighted or hybrid mode	238
Trust Mode	Selects DSCP or CoS priority processing	244
DSCP to DSCP		245
Add	Maps DSCP values in incoming packets to per-hop behavior and drop precedence values for internal priority processing	245
Show	Shows the DSCP to DSCP mapping list	245
CoS to DSCP		248
Add	Maps CoS/CFI values in incoming packets to per-hop behavior and drop precedence values for priority processing	248
Show	Shows the CoS to DSCP mapping list	248
PHB to Queue		241
Add	Maps internal per-hop behavior values to hardware queues	241
Show	Shows the PHB to Queue mapping list	241
DiffServ		251
Configure Class		252
Add	Creates a class map for a type of traffic	252
Show	Shows configured class maps	252
Modify	Modifies the name of a class map	252
Add Rule	Configures the criteria used to classify ingress traffic	252
Show Rule	Shows the traffic classification rules for a class map	252
Configure Policy		256
Add	Creates a policy map to apply to multiple interfaces	256
Show	Shows configured policy maps	256
Modify	Modifies the name of a policy map	256
Add Rule	Sets the boundary parameters used for monitoring inbound traffic, and the action to take for conforming and non-conforming traffic	256
Show Rule	Shows the rules used to enforce bandwidth policing for a policy map	256
Configure Interface	Applies a policy map to an ingress port	265
VoIP	Voice over IP	267
Configure Global	Configures auto-detection of VoIP traffic, sets the Voice VLAN, and VLAN aging time	268

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure OUI		269
Add	Maps the OUI in the source MAC address of ingress packets to the VoIP device manufacturer	269
Show	Shows the OUI telephony list	269
Configure Interface	Configures VoIP traffic settings for ports, including the way in which a port is added to the Voice VLAN, filtering of non-VoIP packets, the method of detecting VoIP traffic, and the priority assigned to the voice traffic	270
Security		273
AAA	Authentication, Authorization and Accounting	274
System Authentication	Configures authentication sequence – local, RADIUS, and TACACS	275
Server		276
Configure Server	Configures RADIUS and TACACS server message exchange settings	276
Configure Group		276
Add	Specifies a group of authentication servers and sets the priority sequence	276
Show	Shows the authentication server groups and priority sequence	276
Accounting	Enables accounting of requested services for billing or security purposes	281
Configure Global	Specifies the interval at which the local accounting service updates information to the accounting server	281
Configure Method		281
Add	Configures accounting for various service types	281
Show	Shows the accounting settings used for various service types	281
Configure Service	Sets the accounting method applied to specific interfaces for 802.1X, CLI command privilege levels for the console port, and for Telnet	281
Show Information		281
Summary	Shows the configured accounting methods, and the methods applied to specific interfaces	281
Statistics	Shows basic accounting information recorded for user sessions	281
Authorization	Enables authorization of requested services	287
Configure Method		287
Add	Configures authorization for various service types	287
Show	Shows the authorization settings used for various service types	287
Configure Service	Sets the authorization method applied used for the console port, and for Telnet	287
Show Information	Shows the configured authorization methods, and the methods applied to specific interfaces	287
User Accounts		290
Add	Configures user names, passwords, and access levels	290

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show	Shows authorized users	290
Modify	Modifies user attributes	290
Web Authentication	Allows authentication and access to the network when 802.1X or Network Access authentication are infeasible or impractical	292
Configure Global	Configures general protocol settings	293
Configure Interface	Enables Web Authentication for individual ports	294
Network Access	MAC address-based network access authentication	295
Configure Global	Enables aging for authenticated MAC addresses, and sets the time period after which a connected MAC address must be reauthenticated	297
Configure Interface		298
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	298
Link Detection	Configures detection of changes in link status, and the response (i.e., send trap or shut down port)	300
Configure MAC Filter		301
Add	Specifies MAC addresses exempt from authentication	301
Show	Shows the list of exempt MAC addresses	301
Show Information	Shows the authenticated MAC address list	303
HTTPS	Secure HTTP	305
Configure Global	Enables HTTPS, and specifies the UDP port to use	305
Copy Certificate	Replaces the default secure-site certificate	306
SSH	Secure Shell	308
Configure Global	Configures SSH server settings	310
Configure Host Key		312
Generate	Generates the host key pair (public and private)	312
Show	Displays RSA and DSA host keys; deletes host keys	312
Configure User Key		313
Copy	Imports user public keys from TFTP server	313
Show	Displays RSA and DSA user keys; deletes user keys	313
ACL	Access Control Lists	315
Configure Time Range	Configures the time to apply an ACL	316
Add	Specifies the name of a time range	316
Show	Shows the name of configured time ranges	316
Add Rule		316
Absolute	Sets exact time or time range	316
Periodic	Sets a recurrent time	316

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show Rule	Shows the time specified by a rule	316
Configure ACL		320
Show TCAM	Shows utilization parameters for TCAM	319
Add	Adds an ACL based on IP or MAC address filtering	320
Show	Shows the name and type of configured ACLs	320
Add Rule	Configures packet filtering based on IP or MAC addresses and other packet attributes	320
Show Rule	Shows the rules specified for an ACL	320
Configure Interface	Binds a port to the specified ACL and time range	
Configure	Binds a port to the specified ACL and time range	332
Add Mirror	Mirrors matching traffic to the specified port	333
Show Mirror	Shows ACLs mirrored to specified port	333
Show Hardware Counters	Shows statistics for ACL hardware counters	335
ARP Inspection		336
Configure General	Enables inspection globally, configures validation of additional address components, and sets the log rate for packet inspection	337
Configure VLAN	Enables ARP inspection on specified VLANs	339
Configure Interface	Sets the trust mode for ports, and sets the rate limit for packet inspection	341
Show Information		342
Show Statistics	Displays statistics on the inspection process	342
Show Log	Shows the inspection log list	343
IP Filter		344
Add	Sets IP addresses of clients allowed management access via the web, SNMP, and Telnet	344
Show	Shows the addresses to be allowed management access	344
Port Security	Configures per port security, including status, response for security breach, and maximum allowed MAC addresses	346
Port Authentication	IEEE 802.1X	348
Configure Global	Enables authentication and EAPOL pass-through	350
Configure Interface	Sets authentication parameters for individual ports	351
Show Statistics	Displays protocol statistics for the selected port	358
DoS Protection	Protects against Denial-of-Service attacks	360

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
IP Source Guard	Filters IP traffic based on static entries in the IP Source Guard table, or dynamic entries in the DHCP Snooping table	363
Port Configuration	Enables IP source guard and selects filter type per port	363
Static Binding		365
Configure ACL Table		365
Add	Adds static addresses to the source guard ACL binding table	365
Show	Shows static addresses in the source guard ACL binding table	365
Configure MAC Table		365
Add	Adds static addresses to the source guard MAC address binding table	365
Show	Shows static addresses in the source guard MAC address binding table	365
Dynamic Binding	Displays the source-guard binding table for a selected interface	368
IPv6 Source Guard	Filters IPv6 traffic based on static entries in the IP Source Guard table, or dynamic entries in the DHCP Snooping table	369
Port Configuration	Enables IPv6 source guard and selects filter type per port	369
Static Binding		371
Add	Adds a static addresses to the source-guard binding table	371
Show	Shows static addresses in the source-guard binding table	371
Dynamic Binding	Displays the source-guard binding table for a selected interface	374
Administration		383
Log		384
System		384
Configure Global	Stores error messages in local memory	384
Show System Logs	Shows logged error messages	384
Remote	Configures the logging of messages to a remote logging process	386
SMTP	Sends an SMTP client message to a participating server	387
LLDP		389
Configure Global	Configures global LLDP timing parameters	389
Configure Interface		391
Configure General	Sets the message transmission mode; enables SNMP notification; and sets the LLDP attributes to advertise	391
Add CA-Type	Specifies the physical location of the device attached to an interface	395
Show Local Device Information		397
General	Displays general information about the local device	397
Port/Trunk	Displays information about each interface	397

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show Remote Device Information		400
Port/Trunk	Displays information about a remote device connected to a port on this switch	400
Port/Trunk Details	Displays detailed information about a remote device connected to this switch	400
Show Device Statistics		409
General	Displays statistics for all connected remote devices	409
Port/Trunk	Displays statistics for remote devices on a selected port or trunk	409
SNMP	Simple Network Management Protocol	411
Configure Global	Enables SNMP agent status, and sets related trap functions	413
Configure Engine		414
Set Engine ID	Sets the SNMP v3 engine ID on this switch	414
Add Remote Engine	Sets the SNMP v3 engine ID for a remote device	415
Show Remote Engine	Shows configured engine ID for remote devices	415
Configure View		417
Add View	Adds an SNMP v3 view of the OID MIB	417
Show View	Shows configured SNMP v3 views	417
Add OID Subtree	Specifies a part of the subtree for the selected view	417
Show OID Subtree	Shows the subtrees assigned to each view	417
Configure Group		419
Add	Adds a group with access policies for assigned users	419
Show	Shows configured groups and access policies	419
Configure User		
Add Community	Configures community strings and access mode	425
Show Community	Shows community strings and access mode	425
Add SNMPv3 Local User	Configures SNMPv3 users on this switch	426
Show SNMPv3 Local User	Shows SNMPv3 users configured on this switch	426
Change SNMPv3 Local User Group	Assign a local user to a new group	426
Add SNMPv3 Remote User	Configures SNMPv3 users from a remote device	429
Show SNMPv3 Remote User	Shows SNMPv3 users set from a remote device	426
Configure Trap		432
Add	Configures trap managers to receive messages on key events that occur on this switch	432
Show	Shows configured trap managers	432
Configure Notify Filter		
Add	Creates an SNMP notification log	436

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show	Shows the configured notification logs	436
Show Statistics	Shows the status of SNMP communications	438
RMON	Remote Monitoring	440
Configure Global		
Add		
Alarm	Sets threshold bounds for a monitored variable	440
Event	Creates a response event for an alarm	443
Show		440
Alarm	Shows all configured alarms	440
Event	Shows all configured events	443
Configure Interface		
Add		
History	Periodically samples statistics on a physical interface	445
Statistics	Enables collection of statistics on a physical interface	448
Show		
History	Shows sampling parameters for each entry in the history group	445
Statistics	Shows sampling parameters for each entry in the statistics group	448
Show Details		
History	Shows sampled data for each entry in the history group	445
Statistics	Shows sampled data for each entry in the history group	448
Cluster		450
Configure Global	Globally enables clustering for the switch; sets Commander status	451
Configure Member	Adds switch Members to the cluster	452
Show Member	Shows cluster switch member; managed switch members	454
ERPS	Ethernet Ring Protection Switching	455
Configure Global	Activates ERPS globally	459
Configure Domain		460
Add	Creates an ERPS ring	460
Show	Shows list of configured ERPS rings, status, and settings	460
Configure Details	Configures ring parameters	460
Configure Operation	Blocks a ring port using Forced Switch or Manual Switch commands	476
CFM	Connectivity Fault Management	480
Configure Global	Configures global settings, including administrative status, cross-check start delay, link trace, and SNMP traps	484

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure Interface	Configures administrative status on an interface	487
Configure MD	Configure Maintenance Domains	487
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	487
Configure Details	Configures the archive hold time and fault notification settings	487
Show	Shows list of configured maintenance domains	487
Configure MA	Configure Maintenance Associations	492
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	492
Configure Details	Configures detailed settings, including continuity check status and interval level, cross-check status, and alarm indication signal parameters	492
Show	Shows list of configured maintenance associations	492
Configure MEP	Configures Maintenance End Points	496
Add	Configures MEPs at the domain boundary to provide management access for each maintenance association	496
Show	Shows list of configured maintenance end points	496
Configure Remote MEP	Configures Remote Maintenance End Points	498
Add	Configures a static list of remote MEPs for comparison against the MEPs learned through continuity check messages	498
Show	Shows list of configured remote maintenance end points	498
Transmit Link Trace	Sends link trace messages to isolate connectivity faults by tracing the path through a network to the designated target node	500
Transmit Loopback	Sends loopback messages to isolate connectivity faults by requesting a target node to echo the message back to the source	502
Transmit Delay Measure	Sends periodic delay-measure requests to a specified MEP within a maintenance association	504
Show Information		
Show Local MEP	Shows the MEPs configured on this device	506
Show Local MEP Details	Displays detailed CFM information about a specified local MEP in the continuity check database	507
Show Local MIP	Shows the MIPs on this device discovered by the CFM protocol	509
Show Remote MEP	Shows MEPs located on other devices which have been discovered through continuity check messages, or statically configured in the MEP database	510
Show Remote MEP Details	Displays detailed CFM information about a specified remote MEP in the continuity check database	511
Show Link Trace Cache	Shows information about link trace operations launched from this device	513
Show Fault Notification Generator	Displays configuration settings for the fault notification generator	514

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show Continuity Check Error	Displays CFM continuity check errors logged on this device	515
OAM	Operation, Administration, and Maintenance	516
Interface	Enables OAM on specified port, sets the mode to active or passive, and enables the reporting of critical events or errored frame events	516
Counters	Displays statistics on OAM PDUs	519
Event Log	Displays the log for recorded link events	520
Remote Interface	Displays information about attached OAM-enabled devices	521
Remote Loopback	Performs a loopback test on the specified port	522
UDLD	UniDirectional Link Detection	525
Configure Global	Configures the message probe interval, detection interval, and recovery interval	526
Configure Interface	Enables UDLD and aggressive mode which reduces the shut-down delay after loss of bidirectional connectivity is detected	527
Show Information	Displays UDLD neighbor information, including neighbor state, expiration time, and protocol intervals	529
IP		613
General		
Routing Interface		
Add Address	Configures an IP interface for a VLAN	613
Configure Interface	Submits a DHCP request for additional information from a VLAN interface configured with a static IPv4 address	617
Show Address	Shows the IP interfaces assigned to a VLAN	613
Ping	Sends ICMP echo request packets to another node on the network	667
Trace Route	Shows the route packets take to the specified destination	668
ARP	Address Resolution Protocol	670
Configure General	Sets the aging time for dynamic entries in the ARP cache	670
Configure Static Address		672
Add	Statically maps a physical address to an IP address	672
Show	Shows the MAC to IP address static table	672
Show Information	Shows entries in the Address Resolution Protocol (ARP) cache	
Dynamic Address	Shows dynamically learned entries in the IP routing table	674
Other Address	Shows internal addresses used by the switch	674
Statistics	Shows statistics on ARP requests sent and received	674
Routing		
Static Routes		675
Add	Configures static routing entries	675

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show	Shows static routing entries	675
Routing Table		
Show Information	Shows all routing entries, including local, static and dynamic routes	677
Configure ECMP Number	Sets the maximum number of equal-cost paths to the same destination that can be installed in the routing table	678
VRRP	Virtual Router Redundancy Protocol	681
Configure Group ID		682
Add	Adds a VRRP group identifier to a VLAN	682
Show	Shows the VRRP group identifier list	682
Add IP Address	Sets a virtual interface address for a VRRP group	682
Show IP Addresses	Shows the virtual interface address assigned to a VRRP group	682
Configure Detail	Configure detailed settings, such as advertisement interval, preemption, priority, and authentication	682
Show Statistics		
Global Statistics	Displays global statistics for VRRP protocol packet errors	688
Group Statistics	Displays statistics for VRRP protocol events and errors on the specified VRRP group and interface	689
IPv6 Configuration		618
Configure Global	Sets an IPv6 default gateway for traffic with no known next hop	619
Configure Interface	Configures IPv6 interface address using auto-configuration or link-local address, and sets related protocol settings	620
Add IPv6 Address	Adds an global unicast, EUI-64, or link-local IPv6 address to an interface	624
Show IPv6 Address	Show the IPv6 addresses assigned to an interface	626
Show IPv6 Neighbor Cache	Displays information in the IPv6 neighbor discovery cache	628
Show Statistics		629
IPv6	Shows statistics about IPv6 traffic	629
ICMPv6	Shows statistics about ICMPv6 messages	629
UDP	Shows statistics about UDP messages	629
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	635
IP Service		637
DNS	Domain Name Service	637
General		637
Configure Global	Enables DNS lookup; defines the default domain name appended to incomplete host names	637

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Add Domain Name	Defines a list of domain names that can be appended to incomplete host names	638
Show Domain Names	Shows the configured domain name list	638
Add Name Server	Specifies IP address of name servers for dynamic lookup	640
Show Name Servers	Shows the name server address list	640
Static Host Table		641
Add	Configures static entries for domain name to address mapping	641
Show	Shows the list of static mapping entries	641
Modify	Modifies the static address mapped to the selected host name	641
Cache	Displays cache entries discovered by designated name servers	642
DHCP	Dynamic Host Configuration Protocol	643
Client	Specifies the DHCP client identifier for an interface	643
Relay	Specifies DHCP relay servers	645
Snooping		375
Configure Global	Enables DHCP snooping globally, MAC-address verification, information option; and sets the information policy	377
Configure VLAN	Enables DHCP snooping on a VLAN	379
Configure Interface	Sets the trust mode for an interface	380
Show Information	Displays the DHCP Snooping binding information	381
Server		646
Configure Global	Enables DHCP service on this switch	647
Configure Excluded Address		648
Add	Adds excluded addresses	648
Show	Shows excluded addresses	648
Configure Pool		649
Add		649
Network	Add address pool for network groups	649
Host	Add address entry for specified host	649
Show	Shows DHCP pool list	649
Modify	Modifies the specified pool entry	649
Show IP Binding	Displays addresses currently bound to DHCP clients	652
UDP Helper		653
General	Enables UDP helper globally on the switch	654

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Forwarding		654
Add	Specifies the UDP destination ports for which broadcast traffic will be forwarded	654
Show	Shows the list of UDP ports to which broadcast traffic will be forwarded	654
Address		655
Add	Specifies the servers to which designated UDP protocol packets are forwarded	655
Show	Shows the servers to which designated UDP protocol packets are forwarded	655
PPPoE Intermediate Agent	Point-to-Point Protocol over Ethernet Intermediate Agent	657
Configure Global	Enables PPPoE IA on the switch, sets access node identifier, sets generic error message	657
Configure Interface	Enables PPPoE IA on an interface, sets trust status, enables vendor tag stripping, sets circuit ID and remote ID	659
Show Statistics	Shows statistics on PPPoE IA protocol messages	660
Multicast		531
IGMP Snooping		533
General	Enables multicast filtering; configures parameters for multicast snooping	535
Multicast Router		539
Add Static Multicast Router	Assigns ports that are attached to a neighboring multicast router	539
Show Static Multicast Router	Displays ports statically configured as attached to a neighboring multicast router	539
Show Current Multicast Router	Displays ports attached to a neighboring multicast router, either through static or dynamic configuration	539
IGMP Member		541
Add Static Member	Statically assigns multicast addresses to the selected VLAN	541
Show Static Member	Shows multicast addresses statically configured on the selected VLAN	541
Interface		544
Configure VLAN	Configures IGMP snooping per VLAN interface	544
Show VLAN Information	Shows IGMP snooping settings per VLAN interface	544
Configure Port	Configures the interface to drop IGMP query packets or all multicast data packets	550
Configure Trunk	Configures the interface to drop IGMP query packets or all multicast data packets	550
Forwarding Entry	Displays the current multicast groups learned through IGMP Snooping	551
Filter		556
Configure General	Enables IGMP filtering for the switch	556

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure Profile		557
Add	Adds IGMP filter profile; and sets access mode	557
Show	Shows configured IGMP filter profiles	557
Add Multicast Group Range	Assigns multicast groups to selected profile	557
Show Multicast Group Range	Shows multicast groups assigned to a profile	557
Configure Interface	Assigns IGMP filter profiles to port interfaces and sets throttling action	559
Statistics		552
Show Query Statistics	Shows statistics for query-related messages	552
Show VLAN Statistics	Shows statistics for protocol messages, number of active groups	552
Show Port Statistics	Shows statistics for protocol messages, number of active groups	552
Show Trunk Statistics	Shows statistics for protocol messages, number of active groups	552
MLD Snooping		561
General	Enables multicast filtering; configures parameters for IPv6 multicast snooping	561
Interface	Configures Immediate Leave status for a VLAN	563
Multicast Router		564
Add Static Multicast Router	Assigns ports that are attached to a neighboring multicast router	564
Show Static Multicast Router	Displays ports statically configured as attached to a neighboring multicast router	564
Show Current Multicast Router	Displays ports attached to a neighboring multicast router, either through static or dynamic configuration	564
MLD Member		566
Add Static Member	Statically assigns multicast addresses to the selected VLAN	566
Show Static Member	Shows multicast addresses statically configured on the selected VLAN	566
Show Current Member	Shows multicast addresses associated with the selected VLAN, either through static or dynamic configuration	566
Group Information	Displays known multicast groups, member ports, the means by which each group was learned, and the corresponding source list	568
IGMP		569
Proxy	Configures IGMP proxy service for multicast routing	570
Interface	Configures Layer 3 IGMP settings for selected VLAN interface	573
Static Group		575
Add	Configures the router to be a static member of a multicast group on the specified VLAN interface	575
Show	Shows multicast group statically assigned to a VLAN interface	575
Group Information		577
Show Information	Shows the current multicast groups learned through IGMP for each VLAN	577

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show Detail	Shows detailed information on each multicast group associated with a VLAN interface	577
Multicast Routing		745
General	Globally enables IPv4 multicast routing	748
Information		749
Show Summary	Shows each multicast route the switch has learned	749
Show Details	Shows additional information for each multicast route the switch has learned, including RP address, upstream router, and downstream interfaces	749
IPv6 Multicast Routing		745
General	Globally enables IPv6 multicast routing	748
Information		749
Show Summary	Shows each multicast route the switch has learned	749
Show Details	Shows additional information for each multicast route the switch has learned, including RP address, upstream router, and downstream interfaces	749
MVR	Multicast VLAN Registration	579
Configure Global	Configures proxy switching and robustness value	581
Configure Domain	Enables MVR for a domain, sets the MVR VLAN, forwarding priority, and upstream source IP	583
Configure Profile		584
Add	Configures multicast stream addresses	584
Show	Shows multicast stream addresses	584
Associate Profile		584
Add	Maps an address profile to a domain	584
Show	Shows addresses profile to domain mapping	584
Configure Interface	Configures MVR interface type and immediate leave mode; also displays MVR operational and active status	587
Configure Static Group Member		589
Add	Statically assigns MVR multicast streams to an interface	589
Show	Shows MVR multicast streams assigned to an interface	589
Show Member	Shows the multicast groups assigned to an MVR VLAN, the source address of the multicast services, and the interfaces with active subscribers	591
Show Statistics		592
Show Query Statistics	Shows statistics for query-related messages	592
Show VLAN Statistics	Shows statistics for protocol messages and number of active groups	592
Show Port Statistics	Shows statistics for protocol messages and number of active groups	592
Show Trunk Statistics	Shows statistics for protocol messages and number of active groups	592

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
MVR6	Multicast VLAN Registration for IPv6	596
Configure Global	Configures proxy switching and robustness value	597
Configure Domain	Enables MVR for a domain, sets the MVR VLAN, forwarding priority, and upstream source IP	599
Configure Profile		600
Add	Configures multicast stream addresses	600
Show	Shows multicast stream addresses	600
Associate Profile		600
Add	Maps an address profile to a domain	600
Show	Shows addresses profile to domain mapping	600
Configure Interface	Configures MVR interface type and immediate leave mode; also displays MVR operational and active status	603
Configure Port	Configures MVR attributes for a port	603
Configure Trunk	Configures MVR attributes for a trunk	603
Configure Static Group Member		605
Add	Statically assigns MVR multicast streams to an interface	605
Show	Shows MVR multicast streams assigned to an interface	605
Show Member	Shows the multicast groups assigned to an MVR VLAN, the source address of the multicast services, and the interfaces with active subscribers	607
Show Statistics		608
Show Query Statistics	Shows statistics for query-related messages	608
Show VLAN Statistics	Shows statistics for protocol messages, number of active groups	608
Show Port Statistics	Shows statistics for protocol messages, number of active groups	608
Show Trunk Statistics	Shows statistics for protocol messages, number of active groups	608
Routing Protocol		
RIP		692
General		693
Configure	Enables or disables RIP, sets the global RIP attributes and timer values	693
Clear Route	Clears the specified route type or network interface from the routing table	696
Network		697
Add	Sets the network interfaces that will use RIP	697
Show	Shows the network interfaces that will use RIP	697
Passive Interface		699
Add	Stops RIP broadcast and multicast messages from being sent on specified network interfaces	699

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show	Shows the configured passive interfaces	699
Neighbor Address		700
Add	Configures the router to directly exchange routing information with a static neighbor	700
Show	Shows adjacent hosts or interfaces configured as a neighboring router	700
Redistribute		701
Add	Imports external routing information from other routing domains (that is, protocols) into the autonomous system	701
Show	Shows the external routing information to be imported from other routing domains	701
Distance		703
Add	Defines an administrative distance for external routes learned from other routing protocols	703
Show	Shows the administrative distances assigned to external routes learned from other routing protocols	703
Interface		704
Add	Configures RIP parameters for each interface, including send and receive versions, authentication, and method of loopback prevention	704
Show	Shows the RIP parameters set for each interface	704
Modify	Modifies RIP parameters for an interface	704
Statistics		
Show Interface Information	Shows RIP settings, and statistics on RIP protocol messages	708
Show Peer Information	Displays information on neighboring RIP routers	709
Reset Statistics	Clears statistics for RIP protocol messages	709
OSPF	Open Shortest Path First (Version 2)	710
Network Area		711
Add	Defines OSPF area address, area ID, and process ID	711
Show	Shows configured areas	711
Show Process	Show configured processes	711
System		714
Configure	Configures the Router ID, global settings, and default information	714
Show	Shows LSA statistics, administrative status, ABR/ASBR, area count, and version number	717
Area		719
Configure Area		719
Add Area	Adds NSSA or stub	719
Show Area	Shows configured NSSA or stub	719

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure NSSA Area	Configures settings for importing routes into or exporting routes out of not-so-stubby areas	720
Configure Stub Area	Configures default cost, and settings for importing routes into a stub	723
Show Information	Shows statistics for each area, including SPF startups, ABR/ASBR count, LSA count, and LSA checksum	725
Area Range		726
Add	Configures route summaries to advertise at an area boundary	726
Show	Shows route summaries advertised at an area boundary	726
Modify	Modifies route summaries advertised at an area boundary	726
Redistribute		728
Add	Redistributes routes from one routing domain to another	728
Show	Shows route types redistributed to another domain	728
Modify	Modifies configuration settings for redistributed routes	728
Summary Address		730
Add	Aggregates routes learned from other protocols for advertising into other autonomous systems	730
Show	Shows configured summary addresses	730
Interface		730
Show	Shows area ID and designated router settings for each interface	732
Configure by VLAN	Configures OSPF protocol settings and authentication for specified VLAN	732
Configure by Address	Configures OSPF protocol settings and authentication for specified interface address	732
Show MD5 Key	Shows MD5 key ID used for each area	732
Virtual Link		737
Add	Configures a virtual link through a transit area to the backbone	737
Show	Shows virtual links, neighbor address, and state	737
Configure Detailed Settings	Configures detailed protocol and authentication settings	737
Show MD5 Key	Shows the MD5 key ID used for each neighbor	737
Information		
LSDB	Shows information about different OSPF Link State Advertisements (LSAs)	740
Neighbor	Shows information about each OSPF neighbor	743
PIM		753
General	Enables PIM globally for the switch	753
Interface	Enables PIM per interface, and sets the mode to dense or sparse	754
Neighbor	Displays information neighboring PIM routers	759

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
SM		
Configure Global	Configures settings for register messages, and use of the SPT	760
BSR Candidate	Configures the switch as a BSR candidate	761
RP Address		763
Add	Sets a static address for an RP and the associated multicast group(s)	763
Show	Shows the static addresses configured for each RP and the associated multicast groups	763
RP Candidate		764
Add	Advertises the switch as an RP candidate to the BSR for the specified multicast groups	764
Show	Shows the multicast groups for which this switch is advertising itself as an RP candidate to the BSR	764
Show Information		
Show BSR Router	Displays information about the BSR	766
Show RP Mapping	Displays the active RPs and associated multicast routing entries	768
PIM6	PIM for IPv6	753
General	Enables PIM globally for the switch	769
Interface	Enables PIM per interface, and sets the mode to dense or sparse	769
Neighbor	Displays information neighboring PIM routers	775
SM		
Configure Global	Configures settings for register messages, and use of the SPT	775
BSR Candidate	Configures the switch as a BSR candidate	777
RP Address		778
Add	Sets a static address for an RP and the associated multicast group(s)	778
Show	Shows the static addresses configured for each RP and the associated multicast groups	778
RP Candidate		780
Add	Advertises the switch as an RP candidate to the BSR for the specified multicast groups	780
Show	Shows the multicast groups for which this switch is advertising itself as an RP candidate to the BSR	780
Show Information		
Show BSR Router	Displays information about the BSR	782
Show RP Mapping	Displays the active RPs and associated multicast routing entries	783

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.
- ◆ [Stacking](#) – Sets master unit, enables stacking on 10G ports, rennumbers the units in the stack.
- ◆ [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 a web interface for configuring system information. The title bar reads "System > General". Below this, there are six rows of configuration fields:

System Description	GTL-2891
System Object ID	1.3.6.1.4.1.22426.10.1.41.101
System Up Time	0 days, 2 hours, 44 minutes, and 4, 43 seconds
System Name	<input type="text"/>
System Location	<input type="text"/>
System Contact	<input type="text"/>

At the bottom right of the form 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 detector is near the back of the unit.
- ◆ **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 displays the 'System > Switch' page in a web interface. It contains two main sections: 'Main Board Information' and 'Management Software Information'. Below these is a 'Temperature List' table.

Main Board Information	
Serial Number	EC0000000000
Number of Ports	28
Hardware Version	R01
Main Power Status	Up
Redundant Power Status	Not present

Management Software Information	
Role	Master
EPLD Version	R00 V0.09
Loader Version	0.0.1.3
Diagnostics Code Version	0.0.0.1
Operation Code Version	1.2.2.19a

Temperature List Total: 1	
Thermal Detector	Temperature (°C)
1	36

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 up to 10240 bytes for Gigabit Ethernet and 10 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.

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 237.](#))
- ◆ **Static Entry Individual Port** – This switch allows static filtering for unicast and multicast addresses. (Refer to [“Setting Static Addresses” on page 195.](#))
- ◆ **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 159.](#))
- ◆ **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

The screenshot displays the 'System > Capability' configuration page. It features a 'General Capability' section with a 'Jumbo Frame' checkbox set to 'Enabled'. Below this is the 'Bridge Extension' section, which contains a list of settings: 'Extended Multicast Filtering Services' (No), 'Traffic Classes' (Enabled), 'Static Entry Individual Port' (Yes), 'VLAN Version Number' (2), 'VLAN Learning' (IVL), 'Local VLAN Capable' (No), 'Configurable PVID Tagging' (Yes), 'Max Supported VLAN Numbers' (4094), 'Max Supported VLAN ID' (4094), and 'GMRP' (Disabled). At the bottom right of the configuration area are 'Apply' and 'Revert' buttons.

System > Capability	
General Capability	
Jumbo Frame	☒ Enabled
Bridge Extension	
Extended Multicast Filtering Services	No
Traffic Classes	Enabled
Static Entry Individual Port	Yes
VLAN Version Number	2
VLAN Learning	IVL
Local VLAN Capable	No
Configurable PVID Tagging	Yes
Max Supported VLAN Numbers	4094
Max Supported VLAN ID	4094
GMRP	Disabled

Apply Revert

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.

Command Usage

- ◆ When logging into an FTP server, the interface prompts for a user name and password configured on the remote server. Note that “Anonymous” is set as the default user name.
- ◆ The reset command will not be accepted during copy operations to flash memory.

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 download.
8. Select a file on the switch to overwrite or specify a new file name.
9. Then click Apply.

Figure 7: Copy Firmware

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. (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' configuration page. The 'Action' dropdown is set to 'Copy'. The 'Copy Type' dropdown is set to 'Running-Config'. The 'Destination File Name' section has a radio button selected for 'startup1.cfg' and an empty text input field. At the bottom right 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.

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

The screenshot shows the 'System > File' configuration page with the 'Action' dropdown set to 'Set Start-Up'. Below the dropdown is a 'File List' table with 4 files. The first file, 'Level1-L3_V1.2.2.19a.bix', is selected with a radio button. At the bottom right are 'Apply' and 'Revert' buttons.

	File Name	File Type	Start-Up	Modify Time	Size (bytes)
☒	Level1-L3_V1.2.2.19a.bix	Operation Code	Y	2000-12-31 01:51:59	18216876
☐	Level1-L3lite_V1.5.2.15.bix	Operation Code	N	2000-12-31 05:52:04	15557428
☐	Factory_Default_Config.cfg	Config File	N	2000-12-31 06:24:43	476
☒	startup1.cfg	Config File	Y	2000-12-31 01:46:29	2087

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: 4					
	File Name	File Type	Start-Up	Modify Time	Size (bytes)
☐	Level1-L3_V1.2.2.19a.bix	Operation Code	Y	2000-12-31 01:51:59	18216876
☐	Level1-L3Re_V1.5.2.15.bix	Operation Code	N	2000-12-31 05:52:04	15557428
☐	Factory_Default_Config.cfg	Config File	N	2000-12-31 06:24:43	476
☐	startup1.cfg	Config File	Y	2000-12-31 01:46:29	2087

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 *Level1-L3.bix* (using upper case and 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 *Level1-L3.BIX* from the server even though *Level1-L3.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, *level1-L3.bix* and *Level1-L3.bix* are considered to be unique files. Thus, if the upgrade file is stored as *Level1-L3.bix* (or even *level1-L3.bix*) on a case-sensitive server, then the switch (requesting *level1-L3.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 *Level1-L3.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-based configuration interface. At the top, there is a breadcrumb 'System > File'. Below it, the 'Action:' dropdown is set to 'Automatic Operation Code Upgrade'. The 'Automatic Opcode Upgrade' checkbox is checked, with the label 'Enabled' next to it. Below this is a text input field for 'Automatic Upgrade Location URL'. A note in blue text states: 'Note: For automatic upgrades, the operation code file name must be set as Level1-L3.bix.'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 1.5.2.15; new version 1.5.2.16
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 displays the 'System > Time' configuration page. At the top, the breadcrumb 'System > Time' is shown. Below it, a 'Step:' dropdown menu is set to '1. Configure General'. The 'Current Time' is displayed as '2014-5-30 9:59:20'. The 'Maintain Type' dropdown menu is set to 'Manual'. Below this, there are input fields for time and date: '9' for Hours, '59' for Minutes, '20' for Seconds, '5' for Month, '30' for Day, and '2014' for Year. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 the 'System > Time' configuration page. At the top, there is a 'Step:' dropdown menu set to '1. Configure General'. Below this, the 'Current Time' is displayed as '2014-5-30 9:59:20'. The 'Maintain Type' is set to 'SNTP' via 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 the 'System > Time' configuration page. At the top, there is a breadcrumb 'System > Time'. Below it is a 'Step:' dropdown menu set to '1. Configure General'. The main content area displays the 'Current Time' as '2014-5-30 9:59:20'. Below this is the 'Maintain Type' dropdown menu, which is set to 'NTP'. Under the 'NTP Configuration' section, there is an 'Authentication Status' checkbox labeled 'Enabled' which is currently unchecked. Below that is the 'Polling Interval' set to '1024 sec'. At the bottom right of the form are two buttons: 'Apply' and 'Revert'.

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 a web interface for configuring system time. The breadcrumb is 'System > Time'. Below it, there are two dropdown menus: 'Step' set to '2. Configure Time Server' and 'Action' set to 'Configure SNTP Server'. Below these are three text input fields labeled 'SNTP Server IP Address 1', 'SNTP Server IP Address 2', and 'SNTP Server IP Address 3'. The first field contains '10.1.0.19', the second contains '137.62.140.80', and the third contains '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's a breadcrumb 'System > Time'. Below it, a 'Step:' dropdown is set to '2. Configure Time Server' and an 'Action:' dropdown is set to 'Add NTP Server'. The 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 a note '(optional)'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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. At the top, there's a breadcrumb 'System > Time'. Below it, a 'Step:' dropdown is set to '2. Configure Time Server' and an 'Action:' dropdown is set to 'Show NTP Server'. Below the dropdowns, it says 'NTP Server List Total: 1'. There is a table with the following data:

	Server IP Address	Version	Authentication Key
☐	192.168.3.20	3	3

At the bottom right, there are 'Delete' and 'Revert' buttons.

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

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 Authentication Key'. The main form area contains two input fields: 'Authentication Key (1-65535)' with the value '3' and 'Key Context (1-32)' with the value 'S1507N122103J068173M'. At the bottom right of the form are two buttons: 'Apply' and 'Revert'.

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



Setting the Time Zone Use the System > Time (Configure Time Zone) 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

System > Time

Step: 3. Configure Time Zone

☐ Predefined Configuration

☒ User Defined Configuration

Direction: After UTC

Name: UTC

Hours (0-13): 0

Minutes (0-59): 0

Note: The maximum value before UTC is 12:00.
The maximum value after UTC is 13:00.

Apply Revert

Configuring Summer Time

Use the Summer Time page to set the system clock forward during the summer months (also known as daylight savings time).

In some countries or regions, clocks are adjusted through the summer months so that afternoons have more daylight and mornings have less. This is known as Summer Time, or Daylight Savings Time (DST). Typically, clocks are adjusted forward one hour at the start of spring and then adjusted backward in autumn.

Parameters

The following parameters are displayed in the web interface:

General Configuration

- ◆ **Summer Time in Effect** – Shows if the system time has been adjusted.
- ◆ **Status** – Shows if summer time is set to take effect during the specified period.
- ◆ **Name** – Name of the time zone while summer time is in effect, usually an acronym. (Range: 1-30 characters)
- ◆ **Mode** – Selects one of the following configuration modes. (The Mode option can only be managed when the Summer Time Status option has been set to enabled for the switch.)

Predefined Mode – Configures the summer time status and settings for the switch using predefined configurations for several major regions of the world. To specify the time corresponding to your local time when summer time is in effect, select the predefined summer-time zone appropriate for your location.

Table 5: Predefined Summer-Time Parameters

Region	Start Time, Day, Week, & Month	End Time, Day, Week, & Month	Rel. Offset
Australia	00:00:00, Sunday, Week 5 of October	23:59:59, Sunday, Week 5 of March	60 min
Europe	00:00:00, Sunday, Week 5 of March	23:59:59, Sunday, Week 5 of October	60 min
New Zealand	00:00:00, Sunday, Week 1 of October	23:59:59, Sunday, Week 3 of March	60 min
USA	02:00:00, Sunday, Week 2 of March	02:00:00, Sunday, Week 1 of November	60 min

Date Mode – Sets the start, end, and offset times of summer time for the switch on a one-time basis. This mode sets the summer-time zone relative to the currently configured time zone. To specify a time corresponding to your local time when summer time is in effect, you must indicate the number of minutes your summer-time zone deviates from your regular time zone.

- ◆ **Offset** – Summer-time offset from the regular time zone, in minutes.
(Range: 1-120 minutes)
- ◆ **From** – Start time for summer-time offset.
- ◆ **To** – End time for summer-time offset.

Recurring Mode – Sets the start, end, and offset times of summer time for the switch on a recurring basis. This mode sets the summer-time zone relative to the currently configured time zone. To specify a time corresponding to your local time when summer time is in effect, you must indicate the number of minutes your summer-time zone deviates from your regular time zone.

- ◆ **Offset** – Summer-time offset from the regular time zone, in minutes.
(Range: 1-120 minutes)
- ◆ **From** – Start time for summer-time offset.
- ◆ **To** – End time for summer-time offset.

Web Interface

To specify summer time settings:

1. Click SNTP, Summer Time.
2. Select one of the configuration modes, configure the relevant attributes, enable summer time status.
3. Click Apply.

Figure 21: Configuring Summer Time

The screenshot shows a web-based configuration interface for a switch. The breadcrumb trail is 'System > Time'. The current step is '4. Configure Summer Time'. The configuration options are as follows:

- Summer Time in Effect:** Set to 'No'.
- Status:** A checkbox labeled 'Enabled' is present.
- Name:** An empty text input field.
- Mode:** A dropdown menu currently showing 'Predefined'.
- Predefined Mode Configuration:** A section header.
- Daylight Savings:** A dropdown menu currently showing 'Australia'.

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

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.
(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 the “password” command in the *CLI Reference Guide*).

Note: Password checking can be enabled or disabled for logging in 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 Console.
2. Specify the connection parameters as required.
3. Click Apply

Figure 22: Console Port Settings

System > Console	
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
<input type="button" value="Apply"/> <input type="button" value="Revert"/>	

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 Telnet connection can only be configured through the CLI (see the “password” command 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 23: Telnet Connection Settings

The screenshot shows the 'System > Telnet' configuration page. It contains several settings with input fields and checkboxes:

- Telnet Status:** A checkbox labeled 'Enabled' is checked.
- TCP Port (1-65535):** An input field containing the value '23'.
- Max Sessions (0-8):** An input field containing the value '8'.
- Login Timeout (10-300):** An input field containing '300' followed by a 'sec' label.
- Exec Timeout (60-65535):** An input field containing '600' followed by a 'sec' label.
- Password Threshold (1-120):** A checkbox is checked, and an input field contains the value '3'.
- Silent Time (1-65535):** An unchecked checkbox is followed by an empty input field and a 'sec' label.

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

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 24: 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 25: Displaying Memory Utilization

System > Memory Status		
Memory Status		
Free Size	138,711,040 bytes	51%
Used Size	129,724,416 bytes	49%
Total	268,435,456 bytes	

Stacking

This section describes the basic functions which enable a properly connected set of switches to function as a single logical entity for management purposes. For information on how to physically connect units into a stack, see the *Hardware Installation Guide*. For detailed information on how stacking is implemented for this type of switch, refer to “Stack Operations” in the *CLI Reference Guide*.

Setting the Master Unit Use the System > Stacking (Configure Master Button) page to configure a unit as the stack master.

Command Usage

- ◆ The switch must be rebooted to activate this command. Note that the configured setting is not affected by changes to the start-up configuration file.
- ◆ Set the front panel 10G ports to stacking mode with the Configure Stacking Button page prior to rebooting the switch.
- ◆ If the stack has not been initialized, the master button must be disabled on all other units in the stack, and those units rebooted.
- ◆ If the stack has been initialized, and this page is used to configure a new stack master, then the master button on the old master unit must be disabled before rebooting the stack.
- ◆ After the newly configured stack master has been rebooted, the front panel unit identifier will be updated on each unit in the stack.
- ◆ The bootup messages on all slave units will be halted when the master unit is rebooted, and configuration through the CLI will be restricted to the master unit.

Parameters

The following parameters are displayed:

- ◆ **Unit** – Shows the stack members according to assigned identifiers.
- ◆ **Master Button** – Enables the specified unit as the stack master.
(Default: Disabled)

Web Interface

To set the stack master:

1. Click System, Stacking.
2. Select Configure Master Button from the Action list.
3. Select one of the stack members as the master unit.

4. Click Apply.

Figure 26: Setting the Stack Master

System > Stacking	
Action: Configure Master Button	
Master Button List Total: 2	
Unit	Master Button
1	☒ Enabled
2	☐ Enabled

Enabling Stacking Ports Use the System > Stacking (Configure Stacking Button) page to enable stacking on the front panel 10G ports.

Command Usage

- ◆ The stacking ports must be enabled on all stack members.
- ◆ Use the Switch Master Button page to specify one unit as the stack master.
- ◆ Every switch in the stack must be rebooted to activate this command. Note that the configured setting is not affected by changes to the start-up configuration file.

Parameters

The following parameters are displayed:

- ◆ **Status** – Enables stacking on the 10G ports. When the configured status is different from the current status, the switch must be rebooted to activate the configured status. (Default: Disabled)
- ◆ **Current Status** – Shows the currently effective status.
- ◆ **Stacking Up Port** – Shows the port which must be connected to next switch up in the stack.
- ◆ **Stacking Down Port** – Shows the port which must be connected to next switch down in the stack.

Web Interface

To enable stacking on the 10G ports:

1. Click System, Stacking.
2. Select Configure Stacking Button from the Action list.
3. Enable stacking on each stack member.
4. Click Apply.

Figure 27: Enabling Stacking on 10G Ports

System > Stacking

Action: Configure Stacking Button

Status	☒ Enabled
Current Status	Enabled
Stacking Up Port	49
Stacking Down Port	50

Note: When the configured status is different from the current status, the configured status takes effect after reboot.

Apply Revert

Renumbering the Stack

If the units are no longer numbered sequentially after several topology changes or failures, use the System > Stacking (Renumber) page to reset the unit numbers. Just remember to save the new configuration settings to a startup configuration file prior to powering off the stack Master.

Command Usage

- ◆ The startup configuration file maps configuration settings to each switch in the stack based on the unit identification number. You should therefore remember to save the current configuration after renumbering the stack.
- ◆ For a line topology, the stack is numbered from top to bottom, with the first unit in the stack designated at unit 1. For a ring topology, the Master unit is taken as the top of the stack and is numbered as unit 1, and all other units are numbered sequentially down through the ring.

Web Interface

To renumber the units in the stack:

1. Click System, then Renumber.
2. Click OK when the confirmation message appears.

Figure 28: Renumbering the Stack

System > Stacking

Action: Renumber

Renumber Click this button to reset the unit numbers.

Message from webpage

System will be renumbered.

OK Cancel

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. (See ["Saving the Running Configuration to a Local File" on page 87](#)).

Parameters

The following parameters are displayed:

System Reload Information

- ◆ Reload Settings – Displays information on the next scheduled reload and selected reload mode as shown in the following example:
"The switch will be rebooted at March 9 12:00:00 2012. Remaining Time: 0 days, 2 hours, 46 minutes, 5 seconds.
Reloading switch regularly time: 12:00 everyday."
- ◆ **Refresh** – Refreshes reload information. Changes made through the console or to system time may need to be refreshed to display the current settings.
- ◆ **Cancel** – Cancels the current settings shown in this field.

System Reload Configuration

- ◆ **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)

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 29: Restarting the Switch (Immediately)

The screenshot shows the 'System > Reset' page. Under 'System Reload Information', it states 'No configured settings for reloading.' with 'Refresh' and 'Cancel' buttons. Under 'System Reload Configuration', the 'Reset Mode' is set to 'Immediately' with a dropdown arrow, and 'Apply' and 'Revert' buttons are present. A modal dialog box titled 'Message from webpage' is displayed in the foreground, asking 'Do you want to reset the switch immediately?' with 'OK' and 'Cancel' buttons.

Figure 30: Restarting the Switch (In)

The screenshot shows the 'System > Reset' page. Under 'System Reload Information', it provides details: '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.', and 'Reloading switch regularity time: 11:20 everyday.' with 'Refresh' and 'Cancel' buttons. Under 'System Reload Configuration', the 'Reset Mode' is set to 'In' with a dropdown arrow. Below this, there are input fields for 'Reload switch in' with '1' in the 'hours' field and '30' in the 'minutes' field. A note states: 'Note: The specified time must be equal to or less than 24 days.' 'Apply' and 'Revert' buttons are at the bottom.

Figure 31: 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.

RefreshCancel

System Reload Configuration:
Reset Mode At
Reload switch at 19/10/2009 (DD/MM/YYYY) 05:00 (HH:MM)
Warning: You have to setup system time first.Otherwise this function won't work.

ApplyRevert

Figure 32: Restarting the Switch (Regularly)

System > Reset

System Reload Information:
No configured settings for reloading.

RefreshCancel

System Reload Configuration:
Reset Mode Regularly
Time 05:30 (HH:MM)
Period
☒ Daily
☐ Weekly Sunday
☐ Monthly 1
Warning: You have to setup system time first.Otherwise this function won't work.

ApplyRevert

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 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.
- ◆ [Cable Test](#) – Performs cable diagnostics on the specified port.
- ◆ [Trunk Configuration](#) – Configures static or dynamic trunks.
- ◆ [Saving Power](#) – Adjusts the power provided to ports based on the length of the cable used to connect to other devices.
- ◆ [Traffic Segmentation](#) – Configures the uplinks and down links to a segmented group of ports.
- ◆ [VLAN Trunking](#) – Configures a tunnel across one or more intermediate switches which pass traffic for VLAN groups to which they do not belong.

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

- ◆ Auto-negotiation must be disabled before you can configure or force a Gigabit RJ-45 interface to use the Speed/Duplex mode or Flow Control options.
- ◆ When using auto-negotiation, the optimal settings will be negotiated between the link partners based on their advertised capabilities. To set the speed, duplex mode, or flow control under auto-negotiation, the required operation modes must be specified in the capabilities list for an interface.
- ◆ The 1000BASE-T standard does not support forced mode. Auto-negotiation should always be used to establish a connection over any 1000BASE-T port or trunk. If not used, the success of the link process cannot be guaranteed when connecting to other types of switches.
- ◆ The Speed/Duplex mode is fixed at and 10Gfull on the 10GBASE SFP+ ports. When auto-negotiation is enabled, the attributes which can be advertised include the speed, duplex mode, and flow control.



Note: Auto-negotiation is not supported for 1000BASE SFP transceivers used in 10G SFP+ Ports 25-28.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28)
- ◆ **Type** – Indicates the port type. (1000BASE-T, 10GBASE SFP+, or 1000BASE SFP (when this transceiver type is used in an SFP+ port).
- ◆ **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/SFP+ ports, or forced/preferred port type for RJ-45/SFP combination ports.

- **None** - Forced transceiver mode is not used for SFP/SFP+ ports. (This is the default setting for RJ-45 ports and SFP/SFP+ ports.)
- **SFP-Forced 1000SFP** - Always uses the SFP/SFP+ port at 1000 Mbps, Full Duplex.
- **SFP-Forced 10GSFP** - Always uses the SFP+ port at 10 Gbps, Full Duplex.
- ◆ **Autonegotiation** (Port Capabilities) – Allows auto-negotiation to be enabled/disabled. When auto-negotiation is enabled, you need to specify the capabilities to be advertised. When auto-negotiation is disabled, you can force the settings for speed, mode, and flow control. The following capabilities are supported.
 - **10h** - Supports 10 Mbps half-duplex operation.
 - **10f** - Supports 10 Mbps full-duplex operation.
 - **100h** - Supports 100 Mbps half-duplex operation.
 - **100f** - Supports 100 Mbps full-duplex operation.
 - **1000f** - Supports 1000 Mbps full-duplex operation.
 - **10Gf** (10G SFP+ ports only) - Supports 10 Gbps full-duplex operation.
 - **FC** - Flow control can eliminate frame loss by “blocking” traffic from end stations or segments connected directly to the switch when its buffers fill. When enabled, back pressure is used for half-duplex operation and IEEE 802.3-2005 (formally IEEE 802.3x) for full-duplex operation.

Default: Autonegotiation enabled on Gigabit and 10 Gigabit ports;
Advertised capabilities for
1000BASE-T – 10half, 10full, 100half, 100full, 1000full
1000BASE-SX/LX/ZX (SFP / SFP+) – 1000full
10GBASE-SR/LR/ER (SFP+) – 10Gfull
- ◆ **Speed/Duplex** – Allows you to manually set the port speed and duplex mode. (i.e., with auto-negotiation disabled)
- ◆ **Flow Control** – Allows automatic or manual selection of flow control. (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 33: Configuring Connections by Port List

Interface > Port > General

Action: Configure by Port List

Port List Total: 28

Port	Type	Name	Admin	Media Type	Autonegotiation	Speed Duplex	Flow Control
1	1000BASE-T	<input type="text"/>	☒ Enabled	None	☒ Enabled ☒ 10h ☒ 100h ☒ 1000f ☒ 10f ☒ 100f ☐ FC	100full	☒ Enabled
2	1000BASE-T	<input type="text"/>	☒ Enabled	None	☒ Enabled ☒ 10h ☒ 100h ☒ 1000f ☒ 10f ☒ 100f ☐ FC	100full	☒ Enabled
3	1000BASE-T	<input type="text"/>	☒ Enabled	None	☒ Enabled ☒ 10h ☒ 100h ☒ 1000f ☒ 10f ☒ 100f ☐ FC	100full	☒ 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 116](#).

Web Interface

To configure port connection parameters:

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

Figure 34: Configuring Connections by Port Range

Interface > Port > General

Action: Configure by Port Range

Port Range (1-28) 1 - 5

Admin ☒ Enabled

Autonegotiation ☒ Enabled

☐ 10h ☐ 100h ☒ 1000f ☒ FC

☒ 10f ☒ 100f ☐ 10Gf

Speed Duplex 10half

Flow Control ☒ Enabled

Apply Revert

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-T, 10GBASE SFP+, or 1000BASE SFP (when this transceiver type is used in an SFP+ port).
- ◆ **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 for SFP/SFP+ ports.
- ◆ **Autonegotiation** – Shows if auto-negotiation is enabled or disabled.
- ◆ **Oper Speed Duplex** – Shows the current speed and duplex mode.
- ◆ **Oper Flow Control** – Shows the flow control type used.

Web Interface

To display port connection parameters:

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

Figure 35: Displaying Port Information

Interface > Port > General

Action: Show Information


Port List Total: 28

Port	Type	Name	Admin	Oper Status	Media Type	Autonegotiation	Oper Speed Duplex	Oper Flow Control
1	1000BASE-T		Enabled	Up	None	Enabled	100full	None
2	1000BASE-T		Enabled	Down	None	Enabled	1000full	None
3	1000BASE-T		Enabled	Down	None	Enabled	1000full	None
4	1000BASE-T		Enabled	Down	None	Enabled	1000full	None
5	1000BASE-T		Enabled	Down	None	Enabled	1000full	None

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 36: 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), or from one or more source ports on remote switches to a destination port on this switch (remote port mirroring as described in [“Configuring Remote Port Mirroring” on page 122](#)).
- ◆ Monitor port speed should match or exceed source port speed, otherwise traffic may be dropped from the monitor port.
- ◆ When mirroring VLAN traffic (see [“Configuring VLAN Mirroring” on page 189](#)) or packets based on a source MAC address (see [“Configuring MAC Address Mirroring” on page 200](#)), the target port cannot be set to the same target ports as that used for port mirroring by this command.
- ◆ When traffic matches the rules for both port mirroring, and for mirroring of VLAN traffic or packets based on a MAC address, the matching packets will not be sent to target port specified for port mirroring.

- ◆ The destination port cannot be a trunk or trunk member port.
- ◆ 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 37: 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 38: Displaying Local Port Mirror Sessions

Interface > Port > Mirror

Action: Show

Mirror Session List Total: 2


	Source (Unit/Port)	Target (Unit/Port)	Type
☐	1 / 7	1 / 8	Both
☐	1 / 9	1 / 10	Rx

Delete Revert

Configuring Remote Port Mirroring

Use the Interface > 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 39: 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 120](#)), 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 162](#)) 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 port¹. 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 port¹, 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

1. Only 802.1Q trunk or hybrid (i.e., general use) ports can be configured as an RSPAN uplink or destination ports – access ports are not allowed (see [“Adding Static Members to VLANs” on page 165](#)).

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)
Only one mirror session is allowed, including both local and remote mirroring. If local mirroring is enabled (see [page 120](#)), then no session can be configured for RSPAN.
- ◆ **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 162](#)).
- ◆ **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 ports¹ 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.
- ◆ **Type** – Specifies the traffic type to be mirrored remotely. (Options: Rx, Tx, Both)

- ◆ **Destination Port** – Specifies the destination port¹ 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.
- ◆ **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 40: Configuring Remote Port Mirroring (Source)

Interface > RSPAN

Session: 1

Operation Status: Up

Switch Role: Source

Remote VLAN: 2

Uplink Port: 4

Source Port Configuration List Total: 28

Source Port	Type
1	Rx
2	Rx
3	None
4	None
5	Tx

Figure 41: Configuring Remote Port Mirroring (Intermediate)

Interface > RSPAN

Session: 1

Operation Status: Up

Switch Role: Intermediate

Remote VLAN: 2

Uplink Port List Total: 28

Port	Uplink
1	☐
2	☐
3	☐
4	☒
5	☐

Figure 42: Configuring Remote Port Mirroring (Destination)

Interface > RSPAN

Session: 1

Operation Status: Up

Switch Role: Destination

Destination Port: 1

Tag: Untagged

Remote VLAN: 2

Uplink Port List Total: 28

Port	Uplink
1	☐
2	☐
3	☐
4	☒
5	☐

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 loading). 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.
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.

Table 6: Port Statistics (Continued)

Parameter	Description
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.
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).

Table 6: Port Statistics (Continued)

Parameter	Description
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.
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 to update the screen.

Figure 43: Showing Port Statistics (Table)

Interface > Port > Statistics

Mode

☒ Interface ☐ Etherlike ☐ RMON ☐ Utilization

Port

1

☐ Auto-refresh

Interface Statistics


Received Octets	1087697	Transmitted Octets	5274231
Received Errors	0	Transmitted Errors	0
Received Unicast Packets	5342	Transmitted Unicast Packets	6891
Received Discarded Packets	0	Transmitted Discarded Packets	0
Received Multicast Packets	3692	Transmitted Multicast Packets	3940
Received Broadcast Packets	145	Transmitted Broadcast Packets	2
Received Unknown Packets	0		

Refresh

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 44: Showing Port Statistics (Chart)



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.
- ◆ **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 45: Displaying Transceiver Data

The screenshot shows the 'Interface > Port > Transceiver' page. At the top, there is a breadcrumb trail. Below it, a 'Port' dropdown menu is set to '11'. The page is divided into two main sections: 'General' and 'DDM Information'. The 'General' section lists various parameters including 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), and Date Code (11-05-25). The 'DDM Information' section lists Temperature (33.85 °C), Vcc (3.28 V), Bias Current (5.50 mA), TX Power (-5.50 dBm), and RX Power (-35.23 dBm).

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	33.85 °C
Vcc	3.28 V
Bias Current	5.50 mA
TX Power	-5.50 dBm
RX Power	-35.23 dBm

Configuring Transceiver Thresholds Use the Interface > Port > Transceiver page to configure thresholds for alarm and warning messages for optical transceivers which support Digital Diagnostic Monitoring (DDM). This page also displays identifying information for supported transceiver types, and operational parameters for transceivers which support DDM.

Parameters

These parameters are displayed:

- ◆ **Port** – Port number. (GTL-2891: SFP+ ports 25-28)
- ◆ **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:** -128.00-128.00 °C
- **Voltage:** 0.00-6.55 Volts
- **Current:** 0.00-131.00 mA
- **Power:** -40.00-8.20 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 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 46: Configuring Transceiver Thresholds

DDM Thresholds

☐ Trap

☐ Auto Mode

	High Alarm	High Warning	Low Warning	Low Alarm
Temperature(°C)	<input type="text" value="75.00"/>	<input type="text" value="70.00"/>	<input type="text" value="0.00"/>	<input type="text" value="-123.00"/>
Voltage(Volts)	<input type="text" value="3.50"/>	<input type="text" value="3.45"/>	<input type="text" value="3.15"/>	<input type="text" value="3.10"/>
Current(mA)	<input type="text" value="100.00"/>	<input type="text" value="90.00"/>	<input type="text" value="7.00"/>	<input type="text" value="6.00"/>
Tx Power(dBm)	<input type="text" value="-9.00"/>	<input type="text" value="-9.50"/>	<input type="text" value="-11.50"/>	<input type="text" value="-12.00"/>
Rx Power(dBm)	<input type="text" value="-3.00"/>	<input type="text" value="-3.50"/>	<input type="text" value="-21.00"/>	<input type="text" value="-21.50"/>

Click this button to restore default DDM thresholds values.

Performing Cable Diagnostics

Use the Interface > Port > Cable Test page to test the cable attached to a port. The cable test will check for any cable faults (short, open, etc.). If a fault is found, the switch reports the length to the fault. Otherwise, it reports the cable length. It can be used to determine the quality of the cable, connectors, and terminations. Problems such as opens, shorts, and cable impedance mismatch can be diagnosed with this test.

Command Usage

- ◆ Cable diagnostics are performed using Digital Signal Processing (DSP) test methods. DSP analyses the cable by sending a pulsed signal into the cable, and then examining the reflection of that pulse.
- ◆ Cable diagnostics can only be performed on twisted-pair media.
- ◆ This cable test is only accurate for Ethernet cables 7 - 100 meters long.
- ◆ The test takes approximately 5 seconds. The switch displays the results of the test immediately upon completion, including common cable failures, as well as the status and approximate length to a fault.
- ◆ Potential conditions which may be listed by the diagnostics include:
 - OK: Correctly terminated pair
 - Open: Open pair, no link partner
 - Short: Shorted pair
 - Not Supported: This message is displayed for any Gigabit Ethernet ports linked up at a speed lower than 1000 Mbps, or for any 10G Ethernet ports.
 - Impedance mismatch: Terminating impedance is not in the reference range.
- ◆ Ports are linked down while running cable diagnostics.

Parameters

These parameters are displayed:

- ◆ **Port** – Switch port identifier.
- ◆ **Type** – Displays media type. (GE – Gigabit Ethernet, Other – SFP/SFP+)
- ◆ **Link Status** – Shows if the port link is up or down.
- ◆ **Test Result** – The results include common cable failures, as well as the status and approximate distance to a fault, or the approximate cable length if no fault is found.

To ensure more accurate measurement of the length to a fault, first disable power-saving mode on the link partner before running cable diagnostics.

For link-down ports, the reported distance to a fault is accurate to within +/- 2 meters. For link-up ports, the accuracy is +/- 10 meters.

- ◆ **Last Updated** – Shows the last time this port was tested.
- ◆ **Test** – Initiates cable test.

Web Interface

To test the cable attached to a port:

1. Click Interface, Port, Cable Test.
2. Click Test for any port to start the cable test.

Figure 47: Performing Cable Tests

Interface > Port > Cable Test

Cable Test Port List Total: 28

Port	Type	Link Status	Test Result (Cable/Fault Distance in Meters)				Last Updated	Action
			Pair A (meters)	Pair B (meters)	Pair C (meters)	Pair D (meters)		
1	GE	Up	OK (0)	OK (0)	Not Supported	Not Supported	2000-12-31 02:11:27	Test
2	GE	Down	Not Tested	Not Tested	Not Tested	Not Tested		Test
3	GE	Down	Not Tested	Not Tested	Not Tested	Not Tested		Test
4	GE	Down	Not Tested	Not Tested	Not Tested	Not Tested		Test
5	GE	Down	Not Tested	Not Tested	Not Tested	Not Tested		Test
6	GE	Down	Not Tested	Not Tested	Not Tested	Not Tested		Test
7	GE	Down	Not Tested	Not Tested	Not Tested	Not Tested		Test
8	GE	Down	Not Tested	Not Tested	Not Tested	Not Tested		Test
9	GE	Down	Not Tested	Not Tested	Not Tested	Not Tested		Test
10	GE	Down	Not Tested	Not Tested	Not Tested	Not Tested		Test

Note: After every test action, wait several seconds and click the refresh button to display test results.

Refresh

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 port 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 16 trunks at a time on the switch, or up to 32 across the stack.

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 trunks before you connect the corresponding network cables between switches to avoid creating a loop.
- ◆ You can create up to 16 trunks on a switch or 32 trunks in the stack, with up to eight 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.
- ◆ Any of the Gigabit ports on the front panel can be trunked together, including ports of different media types.
- ◆ 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. (Range: 1-16)
- ◆ **Member** – The initial trunk member. Use the Add Member page to configure additional members.
 - **Unit** – Unit identifier. (Range: 1)
 - **Port** – Port identifier. (Range: 1-28)

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 (1-26) 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

Interface > Trunk > Static

Step: 1. Configure Trunk Action: Add Member

Trunk 1

Member Unit 1 Port 1

Apply Revert

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 116](#) for a description of the parameters.)
5. Click Apply.

Figure 51: Configuring Connection Parameters for a Static Trunk

Interface > Trunk > Static

Step: 2. Configure General Action: Configure

Static Trunk List Total: 1

Trunk	Type	Name	Admin	Autonegotiation	Speed Duplex	Flow Control	Link Up Down Trap
1	1000BASE-T		Enabled	<div>Enabled</div> <div>10h100h1000f10f100fFC</div>	100full	Enabled	Enabled

Apply Revert

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

Interface > Trunk > Static

Step: 2. Configure General Action: Show Information


Static Trunk List Total: 1

Trunk	Type	Name	Admin	Oper Status	Shutdown Reason	Autonegotiation	Oper Speed Duplex	Oper Flow Control	Link Up Down Trap
1	1000BASE-T		Enabled	Up		Enabled	100full	None	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 described 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: 1-28)
- ◆ **LACP Status** – Enables or disables LACP on a port.

Configure Aggregation Port - Actor/Partner

- ◆ **Port** – Port number. (Range: 1-28)
- ◆ **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)

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 the 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

Interface > Trunk > Dynamic

Step: 2. Configure Aggregation Port Action: Configure

☒ General ☐ Actor ☐ Partner

Port List Total: 28

Port	LACP Status
1	☐ Enabled
2	☒ Enabled
3	☐ Enabled
4	☐ Enabled
5	☐ Enabled

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

Interface > Trunk > Dynamic

Step: 2. Configure Aggregation Port Action: Configure

☐ General ☒ Actor ☐ Partner

Port List Total: 28

Port	Admin Key (0-65535)	System Priority (0-65535)	Port Priority (0-65535)
1	3	32768	32768
2	3	32768	32768
3	1	32768	32768
4	1	32768	32768
5	1	32768	32768

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

Interface > Trunk > Dynamic

Step: 3. Configure Trunk Action: Show Member

Trunk: 1

Member List Total: 2

Member (Unit/Port)
1/1
1/2

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 116](#) for a description of the interface settings.)
5. Click Apply.

Figure 58: Configuring Connection Settings for a Dynamic Trunk

Trunk	Type	Name	Admin	Autonegotiation	Speed Duplex	Flow Control
1	1000BASE-T		☒ Enabled	☒ Enabled ☒ 10h ☒ 100h ☒ 1000f	100full	☒ Enabled

To show 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: Showing Connection Parameters for Dynamic Trunks

Trunk	Type	Name	Admin	Oper Status	Autonegotiation	Oper Speed Duplex	Oper Flow Control
1	1000BASE-T		Enabled	Up	Enabled	1000full	None

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

Interface > Trunk > Dynamic

Step: 2. Configure Aggregation Port Action: Show Information

☒ Counters ☐ Internal ☐ Neighbors

Port 1

Trunk ID 2

Port Counters Information

LACPDUs Sent	29	LACPDUs Received	25
Marker Sent	0	Marker Receive	0
Marker Unknown Pkts	0	Marker Illegal Pkts	0

Refresh

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.
LACPDUs Interval	Number of seconds before invalidating received LACPDU information.

Table 8: LACP Internal Configuration Information (Continued)

Parameter	Description
Admin State, Oper State	<p>Administrative or operational values of the actor's state parameters:</p> <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. ♦ 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

Interface > Trunk > Dynamic

Step: 2. Configure Aggregation Port Action: Show Information

☐ Counters ☒ Internal ☐ Neighbors

Port 1

Trunk ID 2

Port Internal Information

LACP System Priority	32768
LACP Port Priority	32768
Admin Key	3
Oper Key	3
LACPDUs Interval	30 sec
Admin State	Defaulted, Aggregation, Long timeout, LACP-activity
Oper State	Distributing, Collecting, Synchronization, Aggregation, Long timeout, LACP-activity

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 Neighbors.
5. Select a group member from the Port list.

Figure 62: Displaying LACP Port Remote Information

Interface > Trunk > Dynamic

Step: 2. Configure Aggregation Port Action: Show Information

☐ Counters
 ☐ Internal
 ☒ Neighbors

Port 3

Trunk ID 2

Port Neighbors Information

Partner Admin System ID	32768, 00-00-00-00-00-00
Partner Oper System ID	32768, 00-12-CF-61-24-2F
Partner Admin Port Number	3
Partner Oper Port Number	3
Port Admin Priority	32768
Port Oper Priority	32768
Admin Key	0
Oper Key	3
Admin State	Defaulted, Distributing, Collecting, Synchronization, Long timeout
Oper State	Distributing, Collecting, Synchronization, Aggregation, Long timeout, LACP-activity

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.
- ◆ **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



The screenshot shows a web interface for configuring load balancing. The breadcrumb trail at the top reads 'Interface > Trunk > Load Balance'. Below this, there is a label 'Load Balance Mode' followed by a dropdown menu currently showing 'Source and Destination MAC Address'. At the bottom right of the configuration area, there are two buttons: 'Apply' and 'Revert'.

Saving Power

Use the Interface > Green Ethernet page to enable power savings mode on the selected port.

Command Usage

- ◆ IEEE 802.3 defines the Ethernet standard and subsequent power requirements based on cable connections operating at 100 meters. Enabling power saving mode can reduce power used for cable lengths of 60 meters or less, with more significant reduction for cables of 20 meters or less, and continue to ensure signal integrity.
- ◆ The power-saving methods provided by this switch include:
 - Power saving when there is no link partner:
Under normal operation, the switch continuously auto-negotiates to find a link partner, keeping the MAC interface powered up even if no link connection exists. When using power-savings mode, the switch checks for energy on the circuit to determine if there is a link partner. If none is detected, the switch automatically turns off the transmitter, and most of the receive circuitry (entering Sleep Mode). In this mode, the low-power energy-detection circuit continuously checks for energy on the cable. If none is detected, the MAC interface is also powered down to save additional energy. If energy is detected, the switch immediately turns on both the transmitter and receiver functions, and powers up the MAC interface.
 - Power saving when there is a link partner:
Traditional Ethernet connections typically operate with enough power to support at least 100 meters of cable even though average network cable length is shorter. When cable length is shorter, power consumption can be reduced since signal attenuation is proportional to cable length. When power-savings mode is enabled, the switch analyzes cable length to

determine whether or not it can reduce the signal amplitude used on a particular link.



Note: Power savings can only be implemented on Gigabit Ethernet ports when using twisted-pair cabling. Power-savings mode on a active link only works when connection speed is 1 Gbps, and line length is less than 60 meters.

Parameters

These parameters are displayed:

- ◆ **Port** – Power saving mode only applies to the Gigabit Ethernet ports using copper media.
- ◆ **Power Saving Status** – Adjusts the power provided to ports based on the length of the cable used to connect to other devices. Only sufficient power is used to maintain connection requirements. (Default: Enabled on Gigabit Ethernet RJ-45 ports)

Web Interface

To enable power savings:

1. Click Interface, Green Ethernet.
2. Mark the Enabled check box for a port.
3. Click Apply.

Figure 64: Enabling Power Savings

Port	Power Saving Status
21	☒ Enabled
22	☒ Enabled
23	☒ Enabled
24	☒ Enabled
25	☐ Enabled
26	☐ Enabled
27	☐ Enabled
28	☐ Enabled

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 65: Enabling Traffic Segmentation

Interface > Traffic Segmentation

Step: 1. Configure Global

Status ☐ Enabled

Uplink-to-Uplink Mode: Blocking

Apply Revert

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 153](#)).

- ◆ 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: 1-28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-16)

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 66: Configuring Members for Traffic Segmentation

The screenshot shows a web interface titled "Interface > Traffic Segmentation". At the top, there are two dropdown menus: "Step:" set to "2. Configure Session" and "Action:" set to "Add". Below these, the "Session ID (1-4)" is set to "1". The "Direction" is set to "Uplink" via a dropdown. Under the "Interface" section, there are two radio button options: "Port (1-28)" which is selected, and "Trunk (1-26)". The "Port (1-28)" option has a text input field containing "1" followed by a hyphen and an empty field. The "Trunk (1-26)" option has two empty text input fields separated by a hyphen. 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 67: Showing Traffic Segmentation Members

Session ID	Direction	Interface
1	Uplink	Unit 1 / Port 1
1	Downlink	Unit 1 / Port 2

VLAN Trunking


Use the Interface > VLAN Trunking page to allow unknown VLAN groups to pass through the specified interface.

Command Usage

- ◆ Use this feature to configure a tunnel across one or more intermediate switches which pass traffic for VLAN groups to which they do not belong.

The following figure shows VLANs 1 and 2 configured on switches A and B, with VLAN trunking being used to pass traffic for these VLAN groups across switches C, D and E.

Figure 68: Configuring VLAN Trunking



Without VLAN trunking, you would have to configure VLANs 1 and 2 on all intermediate switches – C, D and E; otherwise these switches would drop any frames with unknown VLAN group tags. However, by enabling VLAN trunking on the intermediate switch ports along the path connecting VLANs 1 and 2, you only need to create these VLAN groups in switches A and B. Switches C, D

and E automatically allow frames with VLAN group tags 1 and 2 (groups that are unknown to those switches) to pass through their VLAN trunking ports.

- ◆ VLAN trunking is mutually exclusive with the “access” switchport mode (see [“Adding Static Members to VLANs” on page 165](#)). If VLAN trunking is enabled on an interface, then that interface cannot be set to access mode, and vice versa.
- ◆ To prevent loops from forming in the spanning tree, all unknown VLANs will be bound to a single instance (either STP/RSTP or an MSTP instance, depending on the selected STA mode).
- ◆ If both VLAN trunking and ingress filtering are disabled on an interface, packets with unknown VLAN tags will still be allowed to enter this interface and will be flooded to all other ports where VLAN trunking is enabled. (In other words, VLAN trunking will still be effectively enabled for the unknown VLAN).

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Port** – Port Identifier. (Range: 1-28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-16)
- ◆ **VLAN Trunking Status** – Enables VLAN trunking on the selected interface.

Web Interface

To enable VLAN trunking on a port or trunk:

1. Click Interface, VLAN Trunking.
2. Click Port or Trunk to specify the interface type.
3. Enable VLAN trunking on any of the ports or on a trunk.
4. Click Apply.

Figure 69: Configuring VLAN Trunking

Interface > VLAN Trunking

Interface ☒ Port ☐ Trunk

Port VLAN Trunking List Total: 28 1 2 3

Port	VLAN Trunking Status
1	☒ Enabled
2	☐ Enabled
3	☐ Enabled
4	☐ Enabled
5	☐ Enabled
6	☐ Enabled
7	☐ Enabled
8	☐ Enabled
9	☐ Enabled
10	☐ Enabled

Apply Revert

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.
- ◆ **Protocol VLANs²** – Configures VLAN groups based on specified protocols.
- ◆ **IP Subnet VLANs²** – Maps untagged ingress frames to a specified VLAN if the source address is found in the IP subnet-to-VLAN mapping table.
- ◆ **MAC-based VLANs²** – Maps untagged ingress frames to a specified VLAN if the source MAC address is found in the IP MAC address-to-VLAN mapping table.
- ◆ **VLAN Mirroring** – Mirrors traffic from one or more source VLANs to a target port.
- ◆ **VLAN Translation** – Maps VLAN IDs between the customer and the service provider.

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

2. If a packet matches the rules defined by more than one of these functions, only one of them is applied, with the precedence being MAC-based, IP subnet-based, protocol-based, and then native port-based.

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 and GVRP protocol
- ◆ 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 assign ports on the other VLAN-aware network devices along the path that will carry this traffic to the same VLAN(s), either manually or dynamically using GVRP. 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 70: 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 with GVRP whenever possible to fully automate VLAN registration.

Automatic VLAN Registration – GVRP (GARP VLAN Registration Protocol) defines a system whereby the switch can automatically learn the VLANs to which each end station should be assigned. If an end station (or its network adapter) supports the IEEE 802.1Q VLAN protocol, it can be configured to broadcast a message to your network indicating the VLAN groups it wants to join. When this switch receives these messages, it will automatically place the receiving port in the specified VLANs, and then forward the message to all other ports. When the message arrives at another switch that supports GVRP, it will also place the receiving port in the specified VLANs, and pass the message on to all other ports. VLAN requirements are propagated in this way throughout the network. This allows GVRP-compliant devices to be automatically configured for VLAN groups based solely on end station requests.


To implement GVRP in a network, first add the host devices to the required VLANs (using the operating system or other application software), so that these VLANs can be propagated onto the network. For both the edge switches attached directly to these hosts, and core switches in the network, enable GVRP on the links between these devices. You should also determine security boundaries in the network and

disable GVRP on the boundary ports to prevent advertisements from being propagated, or forbid those ports from joining restricted VLANs.



Note: If you have host devices that do not support GVRP, you should configure static or untagged VLANs for the switch ports connected to these devices (as described in [“Adding Static Members to VLANs” on page 165](#)). But you can still enable GVRP on these edge switches, as well as on the core switches in the network.

Figure 71: Using GVRP



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, set administrative status, or specify Remote VLAN type (see [“Configuring Remote Port Mirroring” on page 122](#)). 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).
VLAN 1 is the default untagged VLAN.
VLAN 4093 is dedicated for [Switch Clustering](#). Configuring this VLAN for other purposes may cause problems in the Clustering operation.
- ◆ **Status** – Enables or disables the specified VLAN.
- ◆ **Remote VLAN** – Reserves this VLAN for RSPAN (see [“Configuring Remote Port Mirroring” on page 122](#)).

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 (see [“Setting the Switch’s IP Address \(IP Version 4\)” on page 613](#)).

Show

- ◆ **VLAN ID** – ID of configured VLAN.
- ◆ **VLAN Name** – Name of the VLAN.
- ◆ **Status** – Operational status of configured VLAN.
- ◆ **Remote VLAN** – Shows if RSPAN is enabled on this VLAN (see [“Configuring Remote Port Mirroring” on page 122](#)).
- ◆ **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. Mark the Status field to configure the VLAN as operational.
5. Specify whether the VLANs are to be used for remote port mirroring.
6. Click Apply.

Figure 72: Creating Static VLANs

The screenshot shows the 'VLAN > Static' configuration page. The 'Action' dropdown is set to 'Add'. The 'VLAN ID (1-4094)' field contains '2'. The 'Status' checkbox is checked and labeled 'Enabled'. The 'Remote VLAN' checkbox is unchecked. At the bottom right are 'Apply' and 'Revert' buttons.

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 or operational 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 73: Modifying Settings for Static VLANs

The screenshot shows the 'VLAN > Static' configuration page. The 'Action' dropdown is set to 'Modify'. The 'VLAN ID (1-4093)' dropdown is set to '2'. The 'VLAN Name' field contains 'R&D'. The 'Status' checkbox is checked and labeled 'Enabled'. The 'L3 Interface' checkbox is checked and labeled 'Enabled'. At the bottom right are 'Apply' and 'Revert' buttons.

To show the configuration settings for VLAN groups:

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

Figure 74: Showing Static VLANs



The screenshot shows the 'VLAN > Static' configuration page. At the top, there is a dropdown menu for 'Action' set to 'Show'. Below this, a section titled 'Static VLAN List' indicates a total of 2 VLANs. A table lists the following VLANs:

	VLAN ID	VLAN Name	Status	Remote VLAN	L3 Interface
☐	1	DefaultVlan	Enabled	Disabled	Enabled
☐	2	R&D	Enabled	Disabled	Enabled

At the bottom of the table, there are 'Delete' and 'Revert' buttons.

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: 1-28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-16)
- ◆ **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.
Access mode is mutually exclusive with VLAN trunking (see [“VLAN Trunking” on page 156](#)). If VLAN trunking is enabled on an interface, then that interface cannot be set to access mode, and vice versa.

- **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.
 - **Forbidden:** Interface is forbidden from automatically joining the VLAN via GVRP. For more information, see “Automatic VLAN Registration” on page 161. This attribute cannot be configured if the specified VLAN does not exist on the switch.

- **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 using Hybrid mode.

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: 1-28)
- ◆ **Trunk Range** – Displays a list of ports. (Range: 1-16)



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. Select a VLAN from the scroll-down list.
4. Set the Interface type to display as Port or Trunk.
5. Modify the settings for any interface as required.
6. Click Apply.

Figure 75: Configuring Static Members by VLAN Index

VLAN > Static

Action: Edit Member by VLAN

VLAN: 1

Interface: ☒ Port ☐ Trunk

Static VLAN Port Member List Total: 28

Port	Mode	PVID	Acceptable Frame Type	Ingress Filtering	Membership Type			
					Tagged	Untagged	Forbidden	None
1	Hybrid	1	All	☐ Enabled	☐	☒	☐	☐
2	Hybrid	1	All	☐ Enabled	☐	☒	☐	☐
3	Hybrid	1	All	☐ Enabled	☐	☒	☐	☐
4	Hybrid	1	All	☐ Enabled	☐	☒	☐	☐
5	Hybrid	1	All	☐ Enabled	☐	☒	☐	☐

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 configure.
4. Modify the settings for any interface as required.
5. Click Apply.

Figure 76: Configuring Static VLAN Members by Interface

VLAN > Static

Action: Edit Member by Interface

Interface: ☒ Port 1 ☐ Trunk 1

Mode: Hybrid

PVID: 1

Acceptable Frame Type: All

Ingress Filtering: ☐ Enabled

Static VLAN Membership List Total: 4

VLAN	Membership Type			
	Tagged	Untagged	Forbidden	None
1	☐	☒	☐	☐
2	☐	☐	☐	☒
3	☐	☐	☐	☒
4	☐	☐	☐	☒

Apply Revert

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 77: Configuring Static VLAN Members by Interface Range

The screenshot shows the 'VLAN > Static' configuration page. At the top, the 'Action' dropdown is set to 'Edit Member by Interface Range'. Below this, the 'Interface' section has radio buttons for 'Port' (selected) and 'Trunk'. The 'Port Range (1-28)' is represented by two input boxes separated by a hyphen. The 'Mode' is a dropdown menu currently set to 'Access'. The 'VLAN ID (1-4094)' is also represented by two input boxes separated by a hyphen. The 'Membership Type' section has four radio buttons: 'Tagged', 'Untagged' (selected), 'Forbidden', and 'None'. At the bottom right, there are 'Apply' and 'Revert' buttons.

Configuring Dynamic VLAN Registration Use the VLAN > Dynamic page to enable GVRP globally on the switch, or to enable GVRP and adjust the protocol timers per interface.

Parameters

These parameters are displayed:

Configure General

- ◆ **GVRP Status** – GVRP defines a way for switches to exchange VLAN information in order to register VLAN members on ports across the network. VLANs are dynamically configured based on join messages issued by host devices and propagated throughout the network. GVRP must be enabled to permit automatic VLAN registration, and to support VLANs which extend beyond the local switch. (Default: Disabled)

Configure Interface

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Port** – Port Identifier. (Range: 1-28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-16)
- ◆ **GVRP Status** – Enables/disables GVRP for the interface. GVRP must be globally enabled for the switch before this setting can take effect (using the Configure General page). When disabled, any GVRP packets received on this port will be discarded and no GVRP registrations will be propagated from other ports. (Default: Disabled)

GVRP cannot be enabled for ports set to Access mode (see [“Adding Static Members to VLANs” on page 165](#)).

- ◆ **GVRP Timers** – Timer settings must follow this rule:
 $2 \times (\text{join timer}) < \text{leave timer} < \text{leaveAll timer}$
 - **Join** – The interval between transmitting requests/queries to participate in a VLAN group. (Range: 20-1000 centiseconds; Default: 20 centiseconds)
 - **Leave** – The interval a port waits before leaving a VLAN group. This time should be set to more than twice the join time. This ensures that after a Leave or LeaveAll message has been issued, the applicants can rejoin before the port actually leaves the group. (Range: 60-3000 centiseconds; Default: 60 centiseconds)
 - **LeaveAll** – The interval between sending out a LeaveAll query message for VLAN group participants and the port leaving the group. This interval should be considerably larger than the Leave Time to minimize the amount of traffic generated by nodes rejoining the group. (Range: 500-18000 centiseconds; Default: 1000 centiseconds)

Show Dynamic VLAN – Show VLAN

VLAN ID – Identifier of a VLAN this switch has joined through GVRP.

VLAN Name – Name of a VLAN this switch has joined through GVRP.

Status – Indicates if this VLAN is currently operational.
(Display Values: Enabled, Disabled)

Show Dynamic VLAN – Show VLAN Member

- ◆ **VLAN** – Identifier of a VLAN this switch has joined through GVRP.
- ◆ **Interface** – Displays a list of ports or trunks which have joined the selected VLAN through GVRP.

Web Interface

To configure GVRP on the switch:

1. Click VLAN, Dynamic.
2. Select Configure General from the Step list.
3. Enable or disable GVRP.
4. Click Apply.

Figure 78: Configuring Global Status of GVRP

The screenshot shows the 'VLAN > Dynamic' configuration page. The 'Step' dropdown is set to '1. Configure General'. The 'GVRP Status' is checked and labeled 'Enabled'. There are 'Apply' and 'Revert' buttons at the bottom right.

To configure GVRP status and timers on a port or trunk:

1. Click VLAN, Dynamic.
2. Select Configure Interface from the Step list.
3. Set the Interface type to display as Port or Trunk.
4. Modify the GVRP status or timers for any interface.
5. Click Apply.

Figure 79: Configuring GVRP for an Interface


The screenshot shows the 'VLAN > Dynamic' configuration page with the 'Step' dropdown set to '2. Configure Interface'. The 'Interface' type is set to 'Port'. Below the interface type, there is a 'Port List' with a total of 28 ports. A table displays the GVRP status and timers for the first five ports.

Port	GVRP Status	GARP Timer (centiseconds)		
		Join (20-1000)	Leave (60-3000)	LeaveAll (500-18000)
1	☐ Enabled	20	60	1000
2	☐ Enabled	20	60	1000
3	☒ Enabled	20	60	1000
4	☐ Enabled	20	60	1000
5	☐ Enabled	20	60	1000

To show the dynamic VLAN joined by this switch:

1. Click VLAN, Dynamic.
2. Select Show Dynamic VLAN from the Step list.
3. Select Show VLAN from the Action list.

Figure 80: Showing Dynamic VLANs Registered on the Switch




The screenshot shows the 'VLAN > Dynamic' configuration page. At the top, there is a 'Step' dropdown menu set to '3. Show Dynamic VLAN' and an 'Action' dropdown menu set to 'Show VLAN'. Below this, the 'Dynamic VLAN List' is displayed with a 'Total: 2' count. The list is a table with three columns: 'VLAN ID', 'VLAN Name', and 'Status'.

VLAN ID	VLAN Name	Status
120		Enabled
150		Enabled

To show the members of a dynamic VLAN:

1. Click VLAN, Dynamic.
2. Select Show Dynamic VLAN from the Step list.
3. Select Show VLAN Members from the Action list.

Figure 81: Showing the Members of a Dynamic VLAN



The screenshot shows the 'VLAN > Dynamic' configuration page. At the top, there is a 'Step' dropdown menu set to '3. Show Dynamic VLAN' and an 'Action' dropdown menu set to 'Show VLAN Member'. Below this, there is a 'VLAN' dropdown menu set to '120'. The 'Dynamic VLAN Member List' is displayed with a 'Total: 10' count. The list is a table with two columns: 'Interface'.

Interface
Unit 1 / Port 1
Unit 1 / Port 4
Unit 1 / Port 10
Unit 1 / Port 16
Unit 1 / Port 19

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 82: 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 178](#). 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 176.](#))
2. Create a Service Provider VLAN, also referred to as an SPVLAN (see [“Configuring VLAN Groups” on page 162.](#))
3. Configure the QinQ tunnel access port to Access mode (see [“Adding an Interface to a QinQ Tunnel” on page 179.](#))
4. Configure the QinQ tunnel access port to join the SPVLAN as an untagged member (see [“Adding Static Members to VLANs” on page 165.](#))
5. Configure the SPVLAN ID as the native VID on the QinQ tunnel access port (see [“Adding Static Members to VLANs” on page 165.](#))
6. Configure the QinQ tunnel uplink port to Uplink mode (see [“Adding an Interface to a QinQ Tunnel” on page 179.](#))
7. Configure the QinQ tunnel uplink port to join the SPVLAN as a tagged member (see [“Adding Static Members to VLANs” on page 165.](#))

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 179](#)). 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 83: Enabling QinQ Tunneling

The screenshot shows a web interface for configuring VLANs. The breadcrumb trail is 'VLAN > Tunnel'. Below this, there is a 'Step:' dropdown menu with '1. Configure Global' selected. The main configuration area has two sections: 'Tunnel Status' with a checkbox labeled 'Enabled' that is checked, and 'Ethernet Type' with a text input field containing '8100'. Below the input field is a note: '(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 172](#).
- ◆ 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: 1-28/52)
- ◆ **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 84: Configuring CVLAN to SPVLAN Mapping Entries

VLAN > Tunnel

Step: 2. Configure Service Action: Add

Port: 1

Customer VLAN ID (1-4094):

Service VLAN ID (1-4094):

Apply 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 85: Showing CVLAN to SPVLAN Mapping Entries

VLAN > Tunnel

Step: 2. Configure Service Action: Show

Port: 1

Tunnel Service Subscriptions List Total: 2

	Customer VLAN ID	Service VLAN ID
☐	1	1
☐	2	200

Delete Revert

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 under "[Enabling QinQ Tunneling on the Switch](#)" 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 176). 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: 1-28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-16)
- ◆ **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 86: Adding an Interface to a QinQ Tunnel

VLAN > Tunnel

Step: 3. Configure Interface

Interface ☒ Port ☐ Trunk

802.1Q Tunnel Port List Total: 28

Port	Mode
1	Uplink
2	Access
3	None
4	None
5	None

Protocol VLANs

The network devices required to support multiple protocols cannot be easily grouped into a common VLAN. This may require non-standard devices to pass traffic between different VLANs in order to encompass all the devices participating in a specific protocol. This kind of configuration deprives users of the basic benefits of VLANs, including security and easy accessibility.

To avoid these problems, you can configure this switch with protocol-based VLANs that divide the physical network into logical VLAN groups for each required protocol. When a frame is received at a port, its VLAN membership can then be determined based on the protocol type being used by the inbound packets.

Command Usage

- ◆ To configure protocol-based VLANs, follow these steps:
 1. First configure VLAN groups for the protocols you want to use (see [“Configuring VLAN Groups” on page 162](#)). Although not mandatory, we suggest configuring a separate VLAN for each major protocol running on your network. Do not add port members at this time.
 2. Create a protocol group for each of the protocols you want to assign to a VLAN using the Configure Protocol (Add) page.
 3. Then map the protocol for each interface to the appropriate VLAN using the Configure Interface (Add) page.
- ◆ When MAC-based, IP subnet-based, or protocol-based VLANs are supported concurrently, priority is applied in this sequence, and then port-based VLANs last.

Configuring Protocol VLAN Groups

Use the VLAN > Protocol (Configure Protocol - Add) page to create protocol groups.

Parameters

These parameters are displayed:

- ◆ **Frame Type** – Choose either Ethernet, RFC 1042, or LLC Other as the frame type used by this protocol.
- ◆ **Protocol Type** – Specifies the protocol type to match. The available options are IP, ARP, RARP and IPv6. If LLC Other is chosen for the Frame Type, the only available Protocol Type is IPX Raw.
- ◆ **Protocol Group ID** – Protocol Group ID assigned to the Protocol VLAN Group. (Range: 1-2147483647)



Note: Traffic which matches IP Protocol Ethernet Frames is mapped to the VLAN (VLAN 1) that has been configured with the switch's administrative IP. IP Protocol Ethernet traffic must not be mapped to another VLAN or you will lose administrative network connectivity to the switch. If lost in this manner, network access can be regained by removing the offending Protocol VLAN rule via the console. Alternately, the switch can be power-cycled, however all unsaved configuration changes will be lost.

Web Interface

To configure a protocol group:

1. Click VLAN, Protocol.
2. Select Configure Protocol from the Step list.
3. Select Add from the Action list.
4. Select an entry from the Frame Type list.
5. Select an entry from the Protocol Type list.
6. Enter an identifier for the protocol group.
7. Click Apply.

Figure 87: Configuring Protocol VLANs

VLAN > Protocol

Step: 1. Configure Protocol Action: Add

Frame Type Ethernet

Protocol Type 08 06 (ARP)

Protocol Group ID (1-2147483647) 1

Apply Revert

To configure a protocol group:

1. Click VLAN, Protocol.
2. Select Configure Protocol from the Step list.
3. Select Show from the Action list.

Figure 88: Displaying Protocol VLANs

VLAN > Protocol

Step: 1. Configure Protocol Action: Show

Protocol to Group Mapping Table Total: 5

	Frame Type	Protocol Type	Protocol Group ID
☐	Ethernet	08 06	1
☐	Ethernet	80 35	2
☐	RFC 1042	08 00	1
☐	RFC 1042	80 35	3
☐	LLC Other	FF FF	5

Delete Revert

Mapping Protocol Groups to Interfaces Use the VLAN > Protocol (Configure Interface - Add) page to map a protocol group to a VLAN for each interface that will participate in the group.

Command Usage

- ◆ When creating a protocol-based VLAN, only assign interfaces using this configuration screen. If you assign interfaces using any of the other VLAN menus such as the VLAN Static table ([page 165](#)), these interfaces will admit traffic of any protocol type into the associated VLAN.
- ◆ When a frame enters a port that has been assigned to a protocol VLAN, it is processed in the following manner:
 - If the frame is tagged, it will be processed according to the standard rules applied to tagged frames.
 - If the frame is untagged and the protocol type matches, the frame is forwarded to the appropriate VLAN.
 - If the frame is untagged but the protocol type does not match, the frame is forwarded to the default VLAN for this interface.

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Port** – Port Identifier. (Range: 1-28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-16)
- ◆ **Protocol Group ID** – Protocol Group ID assigned to the Protocol VLAN Group. (Range: 1-2147483647)
- ◆ **VLAN ID** – VLAN to which matching protocol traffic is forwarded. (Range: 1-4094)

- ◆ **Priority** – The priority assigned to untagged ingress traffic.
(Range: 0-7, where 7 is the highest priority)

Web Interface

To map a protocol group to a VLAN for a port or trunk:

1. Click VLAN, Protocol.
2. Select Configure Interface from the Step list.
3. Select Add from the Action list.
4. Select a port or trunk.
5. Enter the identifier for a protocol group.
6. Enter the corresponding VLAN to which the protocol traffic will be forwarded.
7. Set the priority to assign to untagged ingress frames.
8. Click Apply.

Figure 89: Assigning Interfaces to Protocol VLANs

VLAN > Protocol

Step: 2. Configure Interface Action: Add

Interface ☒ Port 1 ☐ Trunk

Protocol Group ID 1

VLAN ID (1-4094)

Priority (0-7)

Apply Revert

To show the protocol groups mapped to a port or trunk:

1. Click VLAN, Protocol.
2. Select Configure Interface from the Step list.
3. Select Show from the Action list.
4. Select a port or trunk.

Figure 90: Showing the Interface to Protocol Group Mapping

The screenshot shows the 'VLAN > Protocol' configuration page. At the top, there's a breadcrumb 'VLAN > Protocol'. Below it, a 'Step' dropdown is set to '2. Configure interface' and an 'Action' dropdown is set to 'Show'. Under the 'Interface' section, 'Port 1' is selected. Below this is a table titled 'Port To Protocol Group Mapping Table' with a 'Total: 1' count. The table has three columns: 'Protocol Group ID', 'VLAN ID', and 'Priority'. It contains one entry with Protocol Group ID '1', VLAN ID '2', and Priority '0'. At the bottom right of the table are 'Delete' and 'Revert' buttons.

Protocol Group ID	VLAN ID	Priority
1	2	0

Configuring IP Subnet VLANs

Use the VLAN > IP Subnet page to configure IP subnet-based VLANs.

When using port-based classification, all untagged frames received by a port are classified as belonging to the VLAN whose VID (PVID) is associated with that port.

When IP subnet-based VLAN classification is enabled, the source address of untagged ingress frames are checked against the IP subnet-to-VLAN mapping table. If an entry is found for that subnet, these frames are assigned to the VLAN indicated in the entry. If no IP subnet is matched, the untagged frames are classified as belonging to the receiving port's VLAN ID (PVID).

Command Usage

- ◆ Each IP subnet can be mapped to only one VLAN ID. An IP subnet consists of an IP address and a mask. The specified VLAN need not be an existing VLAN.
- ◆ When an untagged frame is received by a port, the source IP address is checked against the IP subnet-to-VLAN mapping table, and if an entry is found, the corresponding VLAN ID is assigned to the frame. If no mapping is found, the PVID of the receiving port is assigned to the frame.
- ◆ The IP subnet cannot be a broadcast or multicast IP address.
- ◆ When MAC-based, IP subnet-based, or protocol-based VLANs are supported concurrently, priority is applied in this sequence, and then port-based VLANs last.

Parameters

These parameters are displayed:

- ◆ **IP Address** – The IP address for a subnet. Valid IP addresses consist of four decimal numbers, 0 to 255, separated by periods.
- ◆ **Subnet Mask** – This mask identifies the host address bits of the IP subnet.


- ◆ **VLAN** – VLAN to which matching IP subnet traffic is forwarded. (Range: 1-4094)
- ◆ **Priority** – The priority assigned to untagged ingress traffic. (Range: 0-7, where 7 is the highest priority; Default: 0)

Web Interface

To map an IP subnet to a VLAN:

1. Click VLAN, IP Subnet.
2. Select Add from the Action list.
3. Enter an address in the IP Address field.
4. Enter a mask in the Subnet Mask field.
5. Enter the identifier in the VLAN field. Note that the specified VLAN need not already be configured.
6. Enter a value to assign to untagged frames in the Priority field.
7. Click Apply.

Figure 91: Configuring IP Subnet VLANs



The screenshot shows a web interface for configuring IP Subnet VLANs. At the top, it says "VLAN > IP Subnet". Below this is an "Action:" dropdown menu with "Add" selected. There are four input fields: "IP Address" with the value "192.168.1.0", "Subnet Mask" with the value "255.255.255.0", "VLAN (1-4094)" with the value "10", and "Priority (0-7)" which is empty. At the bottom right, there are two buttons: "Apply" and "Revert".

To show the configured IP subnet VLANs:

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

Figure 92: Showing IP Subnet VLANs

The screenshot shows a web interface for VLAN configuration. At the top, it says 'VLAN > IP Subnet'. Below that is an 'Action:' dropdown menu set to 'Show'. The main section is titled 'IP Subnet to VLAN Mapping Table' with a 'Total: 1' indicator. It contains a table with the following data:

	IP Address	Subnet Mask	VLAN	Priority
☐	192.168.1.0	255.255.255.0	10	0

Below the table are 'Delete' and 'Revert' buttons.

Configuring MAC-based VLANs

Use the VLAN > MAC-Based page to configure VLAN based on MAC addresses. The MAC-based VLAN feature assigns VLAN IDs to ingress untagged frames according to source MAC addresses.

When MAC-based VLAN classification is enabled, untagged frames received by a port are assigned to the VLAN which is mapped to the frame's source MAC address. When no MAC address is matched, untagged frames are assigned to the receiving port's native VLAN ID (PVID).

Command Usage

- ◆ The MAC-to-VLAN mapping applies to all ports on the switch.
- ◆ Source MAC addresses can be mapped to only one VLAN ID.
- ◆ Configured MAC addresses cannot be broadcast or multicast addresses.
- ◆ When MAC-based, IP subnet-based, or protocol-based VLANs are supported concurrently, priority is applied in this sequence, and then port-based VLANs last.

Parameters

These parameters are displayed:

- ◆ **MAC Address** – A source MAC address which is to be mapped to a specific VLAN. The MAC address must be specified in the format xx-xx-xx-xx-xx-xx.
- ◆ **Mask** – Identifies a range of MAC addresses. (Range: 00-00-00-00-00-00 to ff-ff-ff-ff-ff-ff)

The binary equivalent mask matching the characters in the front of the first non-zero character must all be 1s (e.g., 111, i.e., it cannot be 101 or 001...). A mask for the MAC address: 00-50-6e-00-5f-b1 translated into binary:

MAC: 00000000-01010000-01101110-00000000-01011111-10110001

could be: 11111111-11xxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx

So the mask in hexadecimal for this example could be:

ff-fx-xx-xx-xx-xx/ff-c0-00-00-00-00/ff-e0-00-00-00-00

- ◆ **VLAN** – VLAN to which ingress traffic matching the specified source MAC address is forwarded. (Range: 1-4094)
- ◆ **Priority** – The priority assigned to untagged ingress traffic. (Range: 0-7, where 7 is the highest priority; Default: 0)

Web Interface

To map a MAC address to a VLAN:

1. Click VLAN, MAC-Based.
2. Select Add from the Action list.
3. Enter an address in the MAC Address field, and a mask to indicate a range of addresses if required.
4. Enter an identifier in the VLAN field. Note that the specified VLAN need not already be configured.
5. Enter a value to assign to untagged frames in the Priority field.
6. Click Apply.

Figure 93: Configuring MAC-Based VLANs

VLAN > MAC-Based

Action: Add

MAC Address:

Mask:

VLAN (1-4094):

Priority (0-7):

Apply Revert

To show the MAC addresses mapped to a VLAN:

1. Click VLAN, MAC-Based.
2. Select Show from the Action list.

Figure 94: Showing MAC-Based VLANs

VLAN > MAC-Based			
Action: Show			
MAC-Based VLAN List Total: 1			
	MAC Address	VLAN	Priority
☐	00-AB-CD-11-22-33	10	0
Delete Revert			

Configuring VLAN Mirroring

Use the VLAN > Mirror (Add) page to mirror traffic from one or more source VLANs 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

- ◆ All active ports in a source VLAN are monitored for ingress traffic only.
- ◆ All VLAN mirror sessions must share the same target port, preferably one that is not a member of the source VLAN.
- ◆ When VLAN mirroring and port mirroring are both enabled, they must use the same target port.
- ◆ When VLAN mirroring and port mirroring are both enabled, the target port can receive a mirrored packet twice; once from the source mirror port and again from the source mirrored VLAN.
- ◆ The target port receives traffic from all monitored source VLANs and can become congested. Some mirror traffic may therefore be dropped from the target port.
- ◆ When mirroring VLAN traffic or packets based on a source MAC address (see [“Configuring MAC Address Mirroring” on page 200](#)), the target port cannot be set to the same target ports as that used for port mirroring (see [“Configuring Local Port Mirroring” on page 120](#)).
- ◆ When traffic matches the rules for both port mirroring, and for mirroring of VLAN traffic or packets based on a MAC address, the matching packets will not be sent to target port specified for port mirroring.

Parameters

These parameters are displayed:


- ◆ **Source VLAN** – A VLAN whose traffic will be monitored. (Range: 1-4094)
- ◆ **Target Port** – The destination port that receives the mirrored traffic from the source VLAN. (Range: 1-28/52)

Web Interface

To configure VLAN mirroring:

1. Click VLAN, Mirror.
2. Select Add from the Action list.
3. Select the source VLAN, and select a target port.
4. Click Apply.

Figure 95: Configuring VLAN Mirroring




The screenshot shows the 'VLAN > Mirror' configuration page. At the top, the breadcrumb 'VLAN > Mirror' is displayed. Below it, the 'Action:' dropdown is set to 'Add'. The 'Source VLAN' dropdown is set to '1', and the 'Target Port' dropdown is set to '2'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show the VLANs to be mirrored:

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

Figure 96: Showing the VLANs to Mirror



The screenshot shows the 'VLAN > Mirror' configuration page with the 'Action:' dropdown set to 'Show'. Below the dropdown, there is a table titled 'VLAN Mirror List' with a 'Total: 1' indicator. The table has two columns: 'Source (VLAN)' and 'Target (Unit/Port)'. The first row shows '1' under 'Source (VLAN)' and '1/2' under 'Target (Unit/Port)'. At the bottom right, there are 'Delete' and 'Revert' buttons.

Source (VLAN)	Target (Unit/Port)
1	1/2

Configuring VLAN Translation


Use the VLAN > Translation (Add) page to map VLAN IDs between the customer and service provider for networks that do not support IEEE 802.1Q tunneling.

Command Usage

- ◆ QinQ tunneling uses double tagging to preserve the customer's VLAN tags on traffic crossing the service provider's network. However, if any switch in the path crossing the service provider's network does not support this feature, then the switches directly connected to that device can be configured to swap the customer's VLAN ID with the service provider's VLAN ID for upstream traffic, or the service provider's VLAN ID with the customer's VLAN ID for downstream traffic.

For example, assume that the upstream switch does not support QinQ tunneling. Select Port 1, and set the Old VLAN to 10 and the New VLAN to 100 to map VLAN 10 to VLAN 100 for upstream traffic entering port 1, and VLAN 100 to VLAN 10 for downstream traffic leaving port 1 as shown below.

Figure 97: Configuring VLAN Translation



- ◆ The maximum number of VLAN translation entries is 8 per port, and up to 96 for the system. However, note that configuring a large number of entries may degrade the performance of other processes that also use the TCAM, such as IP Source Guard filter rules, Quality of Service (QoS) processes, QinQ, MAC-based VLANs, VLAN translation, or traps.
- ◆ If VLAN translation is set on an interface, and the same interface is also configured as a QinQ access port on the VLAN > Tunnel (Configure Interface) page, VLAN tag assignments will be determined by the QinQ process, not by VLAN translation.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier.
- ◆ **Old VLAN** – The original VLAN ID. (Range: 1-4094)
- ◆ **New VLAN** – The new VLAN ID. (Range: 1-4094)

Web Interface

To configure VLAN translation:

1. Click VLAN, Translation.
2. Select Add from the Action list.
3. Select a port, and enter the original and new VLAN IDs.
4. Click Apply.

Figure 98: Configuring VLAN Translation

The screenshot shows the 'VLAN > Translation' configuration page. At the top, the breadcrumb 'VLAN > Translation' is displayed. Below it, the 'Action:' dropdown is set to 'Add'. The 'Port' dropdown is set to '1'. There are two input fields: 'Old VLAN (1-4094)' with the value '10' and 'New VLAN (1-4094)' with the value '100'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show the mapping entries for VLANs translation:

1. Click VLAN, Translation.
2. Select Show from the Action list.
3. Select a port.

Figure 99: Showing the Entries for VLAN Translation

The screenshot shows the 'VLAN > Translation' configuration page with the 'Action:' dropdown set to 'Show'. The 'Port' dropdown is set to '1'. Below the form, there is a table titled 'VLAN Translation Table' with a 'Total: 1' indicator. The table has two columns: 'Old VLAN' and 'New VLAN'. The first row shows '10' in the 'Old VLAN' column and '100' in the 'New VLAN' column. At the bottom right, there are 'Delete' and 'Revert' buttons.

	Old VLAN	New VLAN
☐	10	100

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 Address Mirroring](#) – Mirrors traffic matching a specified source address to a target port.
- ◆ [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 195](#)) 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.

- ◆ Also note that MAC address learning cannot be disabled if any of the following conditions exist:
 - 802.1X Port Authentication has been globally enabled on the switch (see [“Configuring 802.1X Global Settings” on page 350](#)).
 - Security Status (see [“Configuring Port Security” on page 346](#)) is enabled on the same interface.

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Port** – Port Identifier. (Range: 1-28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-16)
- ◆ **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 100: Configuring MAC Address Learning

Port	Status
1	☐ Enabled
2	☒ Enabled
3	☒ Enabled
4	☒ Enabled
5	☒ 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 101: Configuring Static MAC Addresses

The screenshot shows the 'MAC Address > Static' configuration page. At the top, there is a breadcrumb 'MAC Address > Static'. Below it, the 'Action:' dropdown is set to 'Add'. The configuration fields include: 'VLAN' set to '1', 'Interface' with radio buttons for 'Port' (selected) and 'Trunk', 'Port' set to '1', 'MAC Address' set to '00-12-cf-94-34-da', and 'Static Status' set to 'Permanent'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show the static addresses in MAC address table:

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

Figure 102: Displaying Static MAC Addresses

The screenshot shows the 'MAC Address > Static' page with the 'Action:' dropdown set to 'Show'. Below the configuration fields, there is a table titled 'Static MAC Address to Interface Mapping Table' with a 'Total: 2' count. The table has columns for checkboxes, MAC Address, VLAN, Interface, Type, and Life Time. Two entries are listed: one for CPU with MAC 00-00-0C-00-00-FD and another for Unit 1 / Port 1 with MAC 00-12-CF-94-34-DA. At the bottom right, there are 'Delete' and 'Revert' buttons.

	MAC Address	VLAN	Interface	Type	Life Time
☐	00-00-0C-00-00-FD	1	CPU	CPU	Delete on Reset
☐	00-12-CF-94-34-DA	1	Unit 1 / Port 1	Config	Permanent

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: 6-7200 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 103: Setting the Address Aging Time

The screenshot shows the 'MAC Address > Dynamic' configuration page. At the top, there is a breadcrumb 'MAC Address > Dynamic'. Below it, an 'Action:' dropdown menu is set to 'Configure Aging'. The main configuration area contains two settings: 'Aging Status' with a checked checkbox labeled 'Enabled', and 'Aging Time (6-672)' with a text input field containing '300' and a unit label 'sec'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

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.
(Values: Learned or Security, the last of which indicates Port Security)
- ◆ **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 104: Displaying the Dynamic MAC Address Table

The screenshot shows the 'MAC Address > Dynamic' web interface. At the top, the 'Action' dropdown is set to 'Show Dynamic MAC'. Below this, the 'Query by:' section has a 'Sort Key' dropdown set to 'MAC Address'. There are three checkboxes: 'MAC Address', 'VLAN', and 'Interface'. The 'VLAN' checkbox is selected, and its value is set to '1'. The 'Interface' checkbox is also selected, and its value is set to 'Port 1'. A 'Query' button is located below the search parameters. Below the search parameters, the 'Dynamic MAC Address List' is displayed with a 'Total: 2' count. The table has five columns: 'MAC Address', 'VLAN', 'Interface', 'Type', and 'Life Time'. The table contains two rows of data.

MAC Address	VLAN	Interface	Type	Life Time
00-E0-29-94-34-64	1	Unit 1 / Port 1	Learn	Delete on Timeout
70-72-CF-32-DD-FF	1	Unit 1 / Port 1	Learn	Delete on Timeout

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 105: Clearing Entries in the Dynamic MAC Address Table



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

Configuring MAC Address Mirroring

Use the MAC Address > Mirror (Add) page to mirror traffic matching a specified source address from any port on the switch 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.

Command Usage

- ◆ When mirroring traffic from a MAC address, ingress traffic with the specified source address entering any port in the switch, other than the target port, will be mirrored to the destination port.
- ◆ All mirror sessions must share the same destination port.
- ◆ Spanning Tree BPDU packets are not mirrored to the target 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 203](#)).
- ◆ When mirroring VLAN traffic (see [“Configuring VLAN Mirroring” on page 189](#)) or packets based on a source MAC address, the target port cannot be set to the same target ports as that used for port mirroring (see [“Configuring Local Port Mirroring” on page 120](#)).
- ◆ When traffic matches the rules for both port mirroring, and for mirroring of VLAN traffic or packets based on a MAC address, the matching packets will not be sent to target port specified for port mirroring.

Parameters

These parameters are displayed:

- ◆ **Source MAC** – MAC address in the form of xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx.
- ◆ **Target Port** – The port that will mirror the traffic from the source port. (Range: 1-28/52)

Web Interface

To mirror packets based on a MAC address:

1. Click MAC Address, Mirror.
2. Select Add from the Action list.
3. Specify the source MAC address and destination port.
4. Click Apply.

Figure 106: Mirroring Packets Based on the Source MAC Address

MAC Address > Mirror

Action: Add

Source MAC: 11-22-33-44-55-66

Target Port: 2

Apply Revert

To show the MAC addresses to be mirrored:

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

Figure 107: Showing the Source MAC Addresses to Mirror

MAC Address > Mirror

Action: Show

MAC Mirror List Total: 1

	Source (MAC)	Target (Unit/Port)
☐	11-22-33-44-55-66	1/2

Delete Revert

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

- ◆ **Port** – Port Identifier. (Range: 1-28/52)
- ◆ **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 108: Issuing MAC Address Traps (Global Configuration)

The screenshot shows the 'MAC Address > MAC Notification' configuration page. The 'Step' dropdown is set to '1. Configure Global'. Below this, there are two settings: 'MAC Notification Traps' with an 'Enabled' checkbox, and 'MAC Notification Trap Interval (1-3600)' with a text input field containing '1' and a 'sec' label. 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 109: Issuing MAC Address Traps (Interface Configuration)

The screenshot shows the 'MAC Address > MAC Notification' configuration page, Step 2: Configure Interface. The 'Step' dropdown is set to '2. Configure Interface'. Below this, there are radio buttons for 'Interface' (selected), 'Port', and 'Trunk'. A 'Port List' section shows 'Total: 28' ports. A table displays the first five ports, each with a 'MAC Notification Trap' checkbox set to 'Enabled'. The table has columns for 'Port' and 'MAC Notification Trap'.

Port	MAC Notification Trap
1	☐ Enabled
2	☐ Enabled
3	☐ Enabled
4	☐ Enabled
5	☐ Enabled

Spanning Tree Algorithm

This chapter describes the following basic topics:

- ◆ [Loopback Detection](#) – Configures detection and response to loopback BPDUs.
- ◆ [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 110: 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 a Internal Spanning Tree (IST) for the Region containing all commonly configured MSTP bridges.

Figure 111: 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 220](#)). 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 112: 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 Loopback Detection

Use the Spanning Tree > Loopback Detection page to configure loopback detection on an interface. When loopback detection is enabled and a port or trunk receives its own BPDU, the detection agent drops the loopback BPDU, sends an SNMP trap, and places the interface in discarding mode. This loopback state can be released manually or automatically. If the interface is configured for automatic loopback release, then the port will only be returned to the forwarding state if one of the following conditions is satisfied:

- ◆ The interface receives any other BPDU except for its own, or;
- ◆ The interface's link status changes to link down and then link up again, or;
- ◆ The interface ceases to receive its own BPDUs in a forward delay interval.



Note: If loopback detection is not enabled and an interface receives its own BPDU, then the interface will drop the loopback BPDU according to IEEE Standard 802.1w-2001 9.3.4 (Note 1).

Note: Loopback detection will not be active if Spanning Tree is disabled on the switch.

Note: When configured for manual release mode, then a link down/up event will not release the port from the discarding state.

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Status** – Enables loopback detection on this interface. (Default: Enabled)
- ◆ **Trap** – Enables SNMP trap notification for loopback events on this interface. (Default: Disabled)
- ◆ **Release Mode** – Configures the interface for automatic or manual loopback release. (Default: Auto)
- ◆ **Release** – Allows an interface to be manually released from discard mode. This is only available if the interface is configured for manual release mode.
- ◆ **Action** – Sets the response for loopback detection to block user traffic or shut down the interface. (Default: Block)
- ◆ **Shutdown Interval** – The duration to shut down the interface. (Range: 60-86400 seconds; Default: 60 seconds)

If an interface is shut down due to a detected loopback, and the release mode is set to “Auto,” the selected interface will be automatically enabled when the shutdown interval has expired.

If an interface is shut down due to a detected loopback, and the release mode is set to “Manual,” the interface can be re-enabled using the Release button.

Web Interface

To configure loopback detection:

1. Click Spanning Tree, Loopback Detection.
2. Click Port or Trunk to display the required interface type.
3. Modify the required loopback detection attributes.
4. Click Apply

Figure 113: Configuring Port Loopback Detection

Spanning Tree > Loopback Detection

Interface ☒ Port ☐ Trunk

Loopback Detection Port List Total: 28

Port	Status	Trap	Release Mode	Release	Action	Shutdown Interval (60-86400 sec)
1	☒ Enabled	☐ Enabled	Auto	Release	Block	
2	☒ Enabled	☐ Enabled	Auto	Release	Block	
3	☒ Enabled	☐ Enabled	Auto	Release	Block	
4	☒ Enabled	☐ Enabled	Auto	Release	Block	
5	☒ Enabled	☒ Enabled	Manual	Release	Shutdown	60

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.

◆ 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,

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

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
- ◆ **BPDU Flooding** – Configures the system to flood BPDUs to all other ports on the switch or just to all other ports in the same VLAN when spanning tree is disabled globally on the switch or disabled on a specific port.
 - To VLAN: Floods BPDUs to all other ports within the receiving port's native VLAN (i.e., as determined by port's PVID). This is the default.
 - To All: Floods BPDUs to all other ports on the switch.

The setting has no effect if BPDU flooding is disabled on a port (see ["Configuring Interface Settings for STA"](#)).

- ◆ **Cisco Prestandard Status** – Configures spanning tree operation to be compatible with Cisco prestandard versions. (Default: Disabled)

Cisco prestandard versions prior to Cisco IOS Release 12.2(25)SEC do not fully follow the IEEE standard, causing some state machine procedures to function incorrectly. This command forces the spanning tree protocol to function in a manner compatible with Cisco prestandard versions.

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; Default: 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)



NOTE: Region Revision and Region Name are both required to uniquely identify an MST region.

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.

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

5. Click Apply

Figure 114: 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

BPDU Flooding To VLAN

Cisco Prestandard Status ☐ Enabled

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 115: 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

BPDU Flooding To VLAN

Cisco Prestandard Status ☐ Enabled

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 116: Configuring Global Settings for STA (MSTP)

The screenshot shows the 'Spanning Tree > STA' configuration page. At the top, there's 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:** 'Enabled' (checked).
- Spanning Tree Type:** 'MSTP' (selected from a dropdown).
- Priority (0-61440, in steps of 4096):** '32768' (text input).
- BPDU Flooding:** 'To VLAN' (selected from a dropdown).
- Cisco Prestandard Status:** 'Enabled' (unchecked).

Advanced:

- Path Cost Method:** 'Long' (selected from a dropdown).
- Transmission Limit (1-10):** '3' (text input).

When the Switch Becomes Root:

- Hello Time (1-10):** '2' sec (text input).
- Maximum Age (6-40):** '20' sec (text input).
- Forward Delay (4-30):** '15' sec (text input).

*Note: 2 * (Hello Time + 1) <= Max Age <= 2 * (Forward Delay - 1)*

MSTP Configuration

- Max Instance Numbers:** '65' (text input).
- Configuration Digest:** '0xAC36177F50283CD4B83821D8AB26DE62' (text input).
- Region Revision (0-65535):** '0' (text input).
- Region Name:** '70 72 cf 83 34 66' (text input).
- Max Hop Count (1-40):** '20' (text input).

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

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.
- ◆ **Configuration 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 117: Displaying Global Settings for STA

Spanning Tree > STA

Step: 1. Configure Global Action: Show Information

Spanning Tree Information

Spanning Tree Status	Enabled	Spanning Tree Type	RSTP
Designated Root	32768.0000E89382A0	Bridge ID	32768.0000E89382A0
Root Port	0	Max Age	20 sec
Root Path Cost	0	Hello Time	2 sec
Configuration Changes	2	Forward Delay	15 sec
Last Topology Change	0 days, 3 hours, 51 minutes, 11 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 207](#)) 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 207](#)). (Default: Enabled)
- ◆ **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


5. Refer to ["Configuring Global Settings for STA" on page 207](#) 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

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 118: 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.)
- ◆ **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 207](#)).

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

- If spanning tree mode is set to STP ([page 207](#)), edge-port mode cannot automatically transition to operational edge-port state using the automatic setting.
- If loopback detection is enabled ([page 205](#)) and a loopback BPDU is detected, the interface cannot function as an edge port until the loopback state is released.
- 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 218](#)).

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 edge port is set to enabled or auto).

- ◆ **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 119: Configuring Interface Settings for STA

Spanning Tree > STA

Step: 2. Configure Interface Action: Configure

Interface ☒ Port ☐ Trunk

Port List Total: 28

Port	Spanning Tree	BPDU Flooding	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 Filter	Migration	Tc Propagate Stop
1	☒ Enabled	☒ Enabled	128	0	Auto	☒ Enabled	Auto	☒ Enabled	☒ Enabled	☒ Enabled	☒ Enabled
2	☒ Enabled	☒ Enabled	128	0	Auto	☒ Enabled	Auto	☒ Enabled	☒ Enabled	☒ Enabled	☒ Enabled
3	☒ Enabled	☒ Enabled	128	0	Auto	☒ Enabled	Auto	☒ Enabled	☒ Enabled	☒ Enabled	☒ Enabled
4	☒ Enabled	☒ Enabled	128	0	Auto	☒ Enabled	Auto	☒ Enabled	☒ Enabled	☒ Enabled	☒ Enabled
5	☒ Enabled	☒ Enabled	128	0	Auto	☒ Enabled	Auto	☒ Enabled	☒ Enabled	☒ Enabled	☒ 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.
- ◆ **BPDU Flooding** – Shows if BPDUs will be flooded to other ports when spanning tree is disabled globally on the switch or disabled on a specific port.
- ◆ **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 213](#).
- ◆ **Oper Edge Port** – This parameter is initialized to the setting for Admin Edge Port in STA Port Configuration on [page 213](#) (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 120: 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 121: Displaying 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: Show Information'. Below these are radio buttons for 'Interface' (selected) and 'Trunk'. The main section is titled 'Spanning Tree Port List' with a 'Total: 28' indicator. A table with 12 columns displays the port settings for five ports. The columns are: Port, Spanning Tree, BPDU Flooding, STA Status, Forward Transitions, Designated Cost, Designated Bridge, Designated Port, Oper Path Cost, Oper Link Type, Oper Edge Port, and Port Role. The data for the five visible ports is as follows:

Port	Spanning Tree	BPDU Flooding	STA Status	Forward Transitions	Designated Cost	Designated Bridge	Designated Port	Oper Path Cost	Oper Link Type	Oper Edge Port	Port Role
1	Enabled	Enabled	Forwarding	3	0	32768.00000C0000FD	128.1	10000	Point-to-Point	Disabled	Designated
2	Enabled	Enabled	Discarding	0	0	32768.00000C0000FD	128.2	10000	Point-to-Point	Disabled	Disabled
3	Enabled	Enabled	Discarding	0	0	32768.00000C0000FD	128.3	10000	Point-to-Point	Disabled	Disabled
4	Enabled	Enabled	Discarding	0	0	32768.00000C0000FD	128.4	10000	Point-to-Point	Disabled	Disabled
5	Enabled	Enabled	Discarding	0	0	32768.00000C0000FD	128.5	10000	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 207) 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 207).

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 122: 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, a 'Step:' dropdown is set to '1. Configure Global' and an 'Action:' dropdown is set to 'Add'. The main form 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 'Apply' and 'Revert' buttons.

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 123: 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' with a subtitle 'Total: 2'. The table has two columns: a checkbox column and an 'MST ID' column. It lists two instances: MST ID 0 and MST ID 1. At the bottom right, there are 'Delete' and 'Revert' buttons.

	MST ID
☐	0
☐	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 MST Instance.
5. Click Apply.

Figure 124: Modifying the Priority for an MST Instance

The screenshot shows the 'Spanning Tree > MSTP' configuration page. The 'Step' dropdown is set to '1. Configure Global' and the 'Action' dropdown is set to 'Modify'. Below this, there is a table titled 'MST Details List' with a total of 2 items. The table has two columns: 'MST ID' and 'Priority (0-61440, in steps of 4096)'. The first row shows MST ID 0 with a priority of 0. The second row shows MST ID 1 with a priority of 32768. At the bottom of the table, there are 'Apply' and 'Revert' buttons.

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

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 212.](#)

Figure 125: Displaying Global Settings for an MST Instance

The screenshot shows the 'Spanning Tree > MSTP' configuration page. The 'Step' dropdown is set to '1. Configure Global' and the 'Action' dropdown is set to 'Show Information'. Below this, there is a dropdown menu for 'MST ID' with '1' selected. The page displays various global settings for the selected MST ID. The settings are organized into two columns.

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 126: Adding a VLAN to an MST Instance

The screenshot shows the 'Spanning Tree > MSTP' configuration page. The 'Step' dropdown is set to '1. Configure Global' and the 'Action' dropdown is set to 'Add Member'. The 'MST ID' dropdown is set to '1'. The 'VLAN ID (1-4094)' text box contains the value '1'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 127: Displaying Members of an MST Instance

The screenshot shows the 'Spanning Tree > MSTP' configuration page with the 'Action' dropdown set to 'Show Member'. The 'MST ID' dropdown is set to '0'. Below the configuration fields, there is a 'Member List' section with a 'Total: 4094' label and a pagination bar showing numbers 1 through 10. A table displays the members of the MST instance.

	VLAN
☐	1
☐	2
☐	3
☐	4
☐	5

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 218](#) 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 207](#)), 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 214](#).

The default path costs are listed in [Table 12 on page 215](#).

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 128: Configuring MSTP Interface Settings

Spanning Tree > STA

Step: 2. Configure Interface Action: Configure

Interface: ☒ Port ☐ Trunk

Port List Total: 28

Port	Spanning Tree	BPDU Flooding	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 Filter	Migration	Tc Propagate Stop
1	☒ Enabled	☒ Enabled	128	0	Auto	☒ Enabled	Auto	☒ Enabled	☒ Enabled	☒ Enabled	☒ Enabled
2	☒ Enabled	☒ Enabled	128	0	Auto	☒ Enabled	Auto	☒ Enabled	☒ Enabled	☒ Enabled	☒ Enabled
3	☒ Enabled	☒ Enabled	128	0	Auto	☒ Enabled	Auto	☒ Enabled	☒ Enabled	☒ Enabled	☒ Enabled
4	☒ Enabled	☒ Enabled	128	0	Auto	☒ Enabled	Auto	☒ Enabled	☒ Enabled	☒ Enabled	☒ Enabled
5	☒ Enabled	☒ Enabled	128	0	Auto	☒ Enabled	Auto	☒ Enabled	☒ Enabled	☒ Enabled	☒ Enabled

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 129: 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: 28

Port	STA Status	Forward Transitions	Designated Cost	Designated Bridge	Designated Port	Oper Path Cost	Oper Link Type	Oper Edge Port	Port Role
1	Forwarding	6	0	0.0.00000C0000FD	128.1	10000	Point-to-Point	Disabled	Designated
2	Discarding	0	0	0.0.00000C0000FD	128.2	10000	Point-to-Point	Disabled	Disabled
3	Discarding	0	0	0.0.00000C0000FD	128.3	10000	Point-to-Point	Disabled	Disabled
4	Discarding	0	0	0.0.00000C0000FD	128.4	10000	Point-to-Point	Disabled	Disabled
5	Discarding	0	0	0.0.00000C0000FD	128.5	10000	Point-to-Point	Disabled	Disabled

Congestion Control

The switch can set the maximum upload or download data transfer rate for any port. It can also control traffic storms by setting a maximum threshold for broadcast traffic or multicast traffic. It can also set bounding thresholds for broadcast and multicast storms which can be used to automatically trigger rate limits or to shut down a port.

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.
- ◆ **Automatic Traffic Control** – Sets thresholds for broadcast and multicast storms which can be used to trigger configured rate limits or to shut down a port.

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 or trunks. 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-T, 10GBASE SFP+, or 1000BASE SFP when this transceiver type is used in an SFP+ port)
- ◆ **Status** – Enables or disables the rate limit. (Default: Disabled)
- ◆ **Rate** – Sets the rate limit level.
(Range: 64 - 1,000,000 kbits per second for Gigabit Ethernet ports;
64 - 10,000,000 kbits per second for 10 Gigabit 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 interface.
4. Set the rate limit for required interfaces.
5. Click Apply.

Figure 130: Configuring Rate Limits



The screenshot shows the 'Traffic > Rate Limit' configuration page. At the top, there's a breadcrumb 'Traffic > Rate Limit'. Below it, the 'Interface' section has radio buttons for 'Port' (selected) and 'Trunk'. A 'Port Rate Limit List' section shows 'Total: 28'. Below this is a table with columns for Port, Type, Input Status, Input Rate, Output Status, and Output Rate. The table lists 5 ports, all of type 1000BASE-T, with 'Enabled' status and a rate of 1000000 kbits/sec for both input and output.

Port	Type	Input		Output	
		Status	Rate (kbits/sec) (64-10000000)	Status	Rate (kbits/sec) (64-10000000)
1	1000BASE-T	☐ Enabled	1000000	☐ Enabled	1000000
2	1000BASE-T	☐ Enabled	1000000	☐ Enabled	1000000
3	1000BASE-T	☐ Enabled	1000000	☐ Enabled	1000000
4	1000BASE-T	☐ Enabled	1000000	☐ Enabled	1000000
5	1000BASE-T	☐ Enabled	1000000	☐ Enabled	1000000

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.
- ◆ Traffic storms can be controlled at the hardware level using Storm Control or at the software level using [Automatic Traffic Control](#) which triggers various

control responses. However, only one of these control types can be applied to a port. Enabling hardware-level storm control on a port will disable automatic storm control on that port.

- ◆ Rate limits set by the storm control function are also used by automatic storm control when the control response is set to rate control on the Auto Traffic Control (Configure Interface) page.
- ◆ 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 the port type. (1000BASE-T, 10GBASE SFP+, or 1000BASE SFP when this transceiver type is used in an SFP+ port)
- ◆ **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 in packets per second. (Range: 500-14881000 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 131: Configuring Storm Control

Traffic > Storm Control							
Interface ☒ Port ☐ Trunk							
Port Storm Control List Total: 28							
Port	Type	Unknown Unicast		Multicast		Broadcast	
		Status	Rate (packets/sec) (500-14881000)	Status	Rate (packets/sec) (500-14880000)	Status	Rate (packets/sec) (500-14881000)
1	1000BASE-T	☐ Enabled	500	☐ Enabled	500	☒ Enabled	500
2	1000BASE-T	☐ Enabled	500	☐ Enabled	500	☒ Enabled	500
3	1000BASE-T	☒ Enabled	6000	☒ Enabled	6000	☒ Enabled	6000
4	1000BASE-T	☐ Enabled	500	☐ Enabled	500	☒ Enabled	500
5	1000BASE-T	☐ Enabled	500	☐ Enabled	500	☒ Enabled	500


Automatic Traffic Control

Use the Traffic > Auto Traffic Control pages to configure bounding thresholds for broadcast and multicast storms which can automatically trigger rate limits or shut down a port.

Command Usage

ATC includes storm control for broadcast or multicast traffic. The control response for either of these traffic types is the same, as shown in the following diagrams.

Figure 132: Storm Control by Limiting the Traffic Rate




The key elements of this diagram are described below:

- ◆ Alarm Fire Threshold – The highest acceptable traffic rate. When ingress traffic exceeds the threshold, ATC sends a Storm Alarm Fire Trap and logs it.
- ◆ When traffic exceeds the alarm fire threshold and the apply timer expires, a traffic control response is applied, and a Traffic Control Apply Trap is sent and logged.

- ◆ Alarm Clear Threshold – The lower threshold beneath which a control response can be automatically terminated after the release timer expires. When ingress traffic falls below this threshold, ATC sends a Storm Alarm Clear Trap and logs it.
- ◆ When traffic falls below the alarm clear threshold after the release timer expires, traffic control (for rate limiting) will be stopped and a Traffic Control Release Trap sent and logged. Note that if the control action has shut down a port, it can only be manually re-enabled using Manual Control Release (see [page 233](#)).
- ◆ The traffic control response of rate limiting can be released automatically or manually. The control response of shutting down a port can only be released manually.

Figure 133: Storm Control by Shutting Down a Port



The key elements of this diagram are the same as that described in the preceding diagram, except that automatic release of the control response is not provided. When traffic control is applied, you must manually re-enable the port.

Functional Limitations

Automatic storm control is a software level control function. Traffic storms can also be controlled at the hardware level using Port Broadcast Control or Port Multicast Control (as described on [page 228](#)). However, only one of these control types can be applied to a port. Enabling automatic storm control on a port will disable hardware-level storm control on that port.

Setting the ATC Timers Use the Traffic > Auto Traffic Control (Configure Global) page to set the time at which to apply the control response after ingress traffic has exceeded the upper threshold, and the time at which to release the control response after ingress traffic has fallen beneath the lower threshold.

Command Usage

- ◆ After the apply timer expires, the settings in the Traffic > Automatic Traffic Control (Configure Interface) page are used to determine if a control action will

be triggered (as configured under the Action field) or a trap message sent (as configured under the Trap Storm Fire field).

- ◆ The release timer only applies to a Rate Control response set in the Action field of the ATC (Interface Configuration) page. When a port has been shut down by a control response, it must be manually re-enabled using the Manual Control Release (see [page 233](#)).

Parameters

These parameters are displayed:

- ◆ **Broadcast Apply Timer** – The interval after the upper threshold has been exceeded at which to apply the control response to broadcast storms.
(Range: 5-300 seconds; Default: 300 seconds)
- ◆ **Broadcast Release Timer** – The time at which to release the control response after ingress traffic has fallen beneath the lower threshold for broadcast storms.
(Range: 5-900 seconds; Default: 900 seconds)
- ◆ **Multicast Apply Timer** – The interval after the upper threshold has been exceeded at which to apply the control response to multicast storms.
(Range: 5-300 seconds; Default: 300 seconds)
- ◆ **Multicast Release Timer** – The time at which to release the control response after ingress traffic has fallen beneath the lower threshold for multicast storms.
(Range: 5-900 seconds; Default: 900 seconds)

Web Interface

To configure the response timers for automatic storm control:

1. Click Traffic, Auto Traffic Control.
2. Select Configure Global from the Step field.
3. Set the apply and release timers for broadcast and multicast storms.
4. Click Apply.

Figure 134: Configuring ATC Timers

The screenshot shows a web interface titled "Traffic > Auto Traffic Control". Below the title is a "Step:" dropdown menu set to "1. Configure Global". The main area contains four rows of configuration fields, each with a label, a text input box, and a unit "sec". The labels are "Broadcast Apply Timer (5-300)", "Broadcast Release Timer (5-900)", "Multicast Apply Timer (5-300)", and "Multicast Release Timer (5-900)". The input boxes contain the values 300, 900, 300, and 900 respectively. At the bottom right of the form are two buttons: "Apply" and "Revert".

Parameter	Value	Unit
Broadcast Apply Timer (5-300)	300	sec
Broadcast Release Timer (5-900)	900	sec
Multicast Apply Timer (5-300)	300	sec
Multicast Release Timer (5-900)	900	sec

Configuring ATC Thresholds and Responses

Use the Traffic > Auto Traffic Control (Configure Interface) page to set the storm control mode (broadcast or multicast), the traffic thresholds, the control response, to automatically release a response of rate limiting, or to send related SNMP trap messages.

Parameters

These parameters are displayed:

- ◆ **Storm Control** – Specifies automatic storm control for broadcast traffic or multicast traffic.
- ◆ **Port** – Port identifier.
- ◆ **State** – Enables automatic traffic control for broadcast or multicast storms. (Default: Disabled)

Automatic storm control is a software level control function. Traffic storms can also be controlled at the hardware level using the [Storm Control](#) menu. However, only one of these control types can be applied to a port. Enabling automatic storm control on a port will disable hardware-level storm control on that port.
- ◆ **Action** – When the Alarm Fire Threshold (upper threshold) is exceeded and the apply timer expires, one of the following control responses will be triggered.
 - **Rate Control** – The rate of ingress traffic is limited to the level set by the Alarm Clear Threshold. Rate limiting is discontinued only after the traffic rate has fallen beneath the Alarm Clear Threshold (lower threshold), and the release timer has expired. (This is the default response.)
 - **Shutdown** – The port is administratively disabled. A port disabled by automatic traffic control can only be manually re-enabled using the Manual Control Release attribute.
- ◆ **Auto Release Control** – Automatically stops a traffic control response of rate limiting when traffic falls below the alarm clear threshold and the release timer expires as illustrated in [Figure 132 on page 230](#). When traffic control stops, the event is logged by the system and a Traffic Release Trap can be sent. (Default: Disabled)

If automatic control release is not enabled and a control response of rate limiting has been triggered, you can manually stop the rate limiting response using the Manual Control Release attribute. If the control response has shut down a port, it can also be re-enabled using Manual Control Release.
- ◆ **Alarm Fire Threshold** – The upper threshold for ingress traffic beyond which a storm control response is triggered after the Apply Timer expires. (Range: 1-255 kilo-packets per second; Default: 128 kpps)

Once the traffic rate exceeds the upper threshold and the Apply Timer expires, a trap message will be sent if configured by the Trap Storm Fire attribute.

- ◆ **Alarm Clear Threshold** – The lower threshold for ingress traffic beneath which a control response for rate limiting will be released after the Release Timer expires, if so configured by the Auto Release Control attribute. (Range: 1-255 kilo-packets per second; Default: 128 kpps)

If rate limiting has been configured as a control response and Auto Control Release is enabled, rate limiting will be discontinued after the traffic rate has fallen beneath the lower threshold, and the Release Timer has expired. Note that if a port has been shut down by a control response, it will not be re-enabled by automatic traffic control. It can only be manually re-enabled using Manual Control Release.

Once the traffic rate falls beneath the lower threshold and the Release Timer expires, a trap message will be sent if configured by the Trap Storm Clear attribute.

- ◆ **Trap Storm Fire** – Sends a trap when traffic exceeds the upper threshold for automatic storm control. (Default: Disabled)
- ◆ **Trap Storm Clear** – Sends a trap when traffic falls beneath the lower threshold after a storm control response has been triggered. (Default: Disabled)
- ◆ **Trap Traffic Apply** – Sends a trap when traffic exceeds the upper threshold for automatic storm control and the apply timer expires. (Default: Disabled)
- ◆ **Trap Traffic Release** – Sends a trap when traffic falls beneath the lower threshold after a storm control response has been triggered and the release timer expires. (Default: Disabled)
- ◆ **Manual Control Release** – Manually releases a control response of rate-limiting or port shutdown any time after the specified action has been triggered.

Web Interface

To configure the response timers for automatic storm control:

1. Click Traffic, Auto Traffic Control.
2. Select Configure Interface from the Step field.
3. Enable or disable ATC as required, set the control response, specify whether or not to automatically release the control response of rate limiting, set the upper and lower thresholds, and specify which trap messages to send.
4. Click Apply.

Figure 135: Configuring ATC Interface Attributes

Traffic > Auto Traffic Control

Step: 2. Configure Interface

Storm Control ☒ Broadcast ☐ Multicast

Auto Traffic Control Broadcast List Total: 28

Port	State	Action	Auto Release Control	Alarm Fire Threshold (kpps) (1-255)	Alarm Clear Threshold (kpps) (1-255)	Trap Storm Fire	Trap Storm Clear	Trap Traffic Apply	Trap Traffic Release	Manual Control Release
1	☐ Enabled	Rate Control	☐ Enabled	128	128	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	Release
2	☐ Enabled	Rate Control	☐ Enabled	128	128	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	Release
3	☐ Enabled	Rate Control	☐ Enabled	128	128	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	Release
4	☐ Enabled	Rate Control	☐ Enabled	128	128	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	Release
5	☐ Enabled	Rate Control	☐ Enabled	128	128	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	Release

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 136: Setting the Default Port Priority

Port	CoS (0-7)
1	0
2	0
3	5
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 service is used for the high priority queues and weighted service for the remaining queues. 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 one of the queuing modes that use a combination of strict and weighted queuing.

- ◆ The specified queue mode applies to all interfaces.

Parameters

These parameters are displayed:

- ◆ **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 priority on the high-priority queues and WRR on the remaining queues.
- ◆ **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 service is used for the high priority queues and weighted service for the remaining queues. 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: 1-255; 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. Set the queue mode.
3. If the weighted queue mode is selected, the queue weight can be modified if required.
4. If the queue mode that uses a combination of strict and weighted queueing is selected, the queues which are serviced first must be specified by enabling strict mode parameter in the table.
5. Click Apply.

Figure 137: Setting the Queue Mode (Strict)

The screenshot shows the 'Traffic > Priority > Queue' configuration page. The 'Queue Mode' dropdown menu is set to 'Strict'. Below the dropdown are 'Apply' and 'Revert' buttons.

Figure 138: Setting the Queue Mode (WRR)

The screenshot shows the 'Traffic > Priority > Queue' configuration page. The 'Queue Mode' dropdown menu is set to 'WRR'. Below the dropdown is a 'Queue Setting Table' with a 'Total: 8' label. The table has two columns: 'Queue ID' and 'Weight (1-255)'. The table contains 8 rows with Queue IDs from 0 to 7 and corresponding weights from 1 to 14. Below the table are 'Apply' and 'Revert' buttons.

Queue ID	Weight (1-255)
0	1
1	2
2	4
3	6
4	8
5	10
6	12
7	14

Figure 139: Setting the Queue Mode (Strict and WRR)

Traffic > Priority > Queue

Queue Mode: Strict and WRR

Queue Setting Table Total: 8

Queue ID	Strict Mode	Weight (1-255)
0	Disabled	1
1	Disabled	2
2	Disabled	4
3	Disabled	6
4	Disabled	8
5	Disabled	10
6	Disabled	12
7	Disabled	14

Apply Revert

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 248](#)).

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 248](#)).

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:

- ◆ **Port** – Specifies a port.
- ◆ **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. Select a port.
4. 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.
5. Click Apply.

Figure 140: Mapping CoS Values to Egress Queues

Traffic > Priority > PHB to Queue

Action: Configure

Port: 1

PHB (0-7):

Queue (0-7):

Apply Revert

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 a port.

Figure 141: Showing CoS Values to Egress Queue Mapping

Traffic > Priority > PHB to Queue

Action: Show

Port: 1

PHB to Queue Mapping List Total: 8

	PHB	Queue
☐	0	2
☐	1	0
☐	2	1
☐	3	3
☐	4	4
☐	5	5
☐	6	6
☐	7	7

Default Revert

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 DSCP or CoS

The switch allows a choice between using 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 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 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 237](#)) 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 237](#)) is used for priority processing.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28)
- ◆ **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.

Web Interface

To configure the trust mode:

1. Click Traffic, Priority, Trust Mode.
2. Set the trust mode for any port.
3. Click Apply.

Figure 142: 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 244](#)), 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 143: Configuring DSCP to DSCP Internal Mapping

Traffic > Priority > DSCP to DSCP

Action: Configure

Port: 1

DSCP (0-63): 1

PHB (0-7): 3

Drop Precedence: 1: Red

Apply Revert

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 144: Showing DSCP to DSCP Internal Mapping

Traffic > Priority > DSCP to DSCP

Action: Show

Port: 1

DSCP to DSCP Mapping List Max: 64 Total: 64

DSCP	PHB	Drop Precedence
0	0	0
1	0	1
2	0	0
3	0	3
4	0	0
5	0	1
6	0	0
7	0	3
8	1	0
9	1	1

Mapping CoS Priorities to Internal DSCP Values

Use the Traffic > Priority > CoS to DSCP page to map 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 248](#).
- ◆ 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.

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 in 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)	(0,0)
1		(1,0)	(1,0)
2		(2,0)	(2,0)
3		(3,0)	(3,0)
4		(4,0)	(4,0)
5		(5,0)	(5,0)
6		(6,0)	(6,0)
7		(7,0)	(7,0)

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 145: Configuring CoS to DSCP Internal Mapping

The screenshot shows a web interface titled "Traffic > Priority > CoS to DSCP". Below the title bar, there is an "Action:" dropdown menu set to "Configure". The main configuration area contains several input fields: "Port" with a dropdown menu showing "1", "CoS (0-7)" with a text input field containing "0", "CFI (0-1)" with a text input field containing "1", "PHB (0-7)" with a text input field containing "0", and "Drop Precedence" with a dropdown menu showing "0: Green". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

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 146: Showing CoS to DSCP Internal Mapping



Traffic > Priority > CoS to DSCP

Action:

Port:

CoS to DSCP Mapping List Total: 16

☐	CoS	CFI	PHB	Drop Precedence
☐	0	0	0	0
☐	0	1	0	0
☐	1	0	1	0
☐	1	1	1	0
☐	2	0	2	0
☐	2	1	2	0
☐	3	0	3	0
☐	3	1	3	0
☐	4	0	4	0
☐	4	1	4	0

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, CoS values, or source ports. 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 256](#)).

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, a CoS value, or a source port.
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 200 classes can be included in a policy map.

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 256](#)) to create a service policy ([page 265](#)) 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** – The criteria specified by the match command.
 - **Match All** – Match all conditions within a class map.

- **Match Any** – Match any condition within a class map.

◆ **Description** – A brief description of a class map. (Range: 1-64 characters)

Add Rule

◆ **Class Name** – Name of the class map.

◆ **Type** – The criteria specified by the match command. (This field is set on the Add page.)

◆ **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)

◆ **IPv6 DSCP** – A DSCP value contained in an IPv6 packet. (Range: 0-63)

◆ **VLAN ID** – A VLAN. (Range:1-4094)

◆ **CoS** – A CoS value. (Range: 0-7)

◆ **Source Port** – A source port. (Range: 1-28)

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. Set the Type to match any rule or all rules.
6. Enter a description.
7. Click Add.

Figure 147: Configuring a Class Map

The screenshot shows the 'Traffic > DiffServ' configuration page. At the top, there is a breadcrumb trail. Below it, the 'Step' dropdown is set to '1. Configure Class' and the 'Action' dropdown is set to 'Add'. The main form contains three fields: 'Class Name' with the value 'rd-class', 'Type' with the value 'Match Any', and 'Description' with the value 'class for software group'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 148: Showing Class Maps

The screenshot shows the 'Traffic > DiffServ' configuration page with the 'Step' dropdown set to '1. Configure Class' and the 'Action' dropdown set to 'Show'. Below the dropdowns, there is a 'Class List' section with a 'Total: 1' indicator. The list contains one entry with the following details:

	Class Name	Type	Description
☐	rd-class	Match Any	class for software group

At the bottom right, there are 'Delete' and 'Revert' buttons.

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, DSCP or IP Precedence value, VLAN, CoS value, or source port. You can specify up to 16 items to match when assigning ingress traffic to a class map.
6. Click Apply.

Figure 149: Adding Rules to a Class Map

Traffic > DiffServ

Step: 1. Configure Class Action: Add Rule

Class Name: rd-class

Type: Match Any

Rule:

☐ ACL

☒ IP DSCP (0-63)

☐ IP Precedence (0-7)

☐ IPv6 DSCP (0-63)

☐ VLAN ID (1-4094)

☐ CoS (0-7)

☐ Source Port

Apply Revert

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 150: Showing the Rules for a Class Map

Traffic > DiffServ

Step: 1. Configure Class Action: Show Rule

Class Name: rd-class

Type: Match Any

Rule List Total: 2

	Rule
☐	IP DSCP 3
☐	IP Precedence 3

Delete Revert

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 252](#)), modify service tagging, and enforce bandwidth policing. A policy map can then be bound by a service policy to one or more interfaces ([page 265](#)).

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 512 class statements that can be applied to the same interface ([page 265](#)). 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 265](#)) 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 up to 200 class maps.
- ◆ **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 248](#)).
 - **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 16, “Default Mapping of DSCP Values to Internal PHB/Drop Values,” on page 246](#)).
 - **Set IP DSCP** – Configures the service provided to ingress traffic by setting an IP DSCP value for a matching packet (as specified in rule settings for a class map). (Range: 0-63)
- ◆ **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 rate in kilobits per second. (Range: 0-10000000 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 in bytes. (Range: 64-16000000 at a granularity of 4k bytes)
The burst size cannot exceed 16 Mbytes.
- **Conform** – Specifies that traffic conforming to the maximum committed rate (CIR) and committed burst size (BC) will be transmitted without any change to the DSCP service level.
 - **Transmit** – Transmits in-conformance traffic without any change to the DSCP service level.
- **Violate** – Specifies whether the traffic that exceeds the maximum rate (CIR) 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 rate in kilobits per second. (Range: 0-10000000 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 in bytes.
(Range: 64-16000000 at a granularity of 4k bytes)
The burst size cannot exceed 16 Mbytes.
- **Excess Burst Size (BE)** – Burst in excess of committed burst size.
(Range: 64-16000000 at a granularity of 4k bytes)
The burst size cannot exceed 16 Mbytes.
- **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.
 - **Transmit** – Transmits in-conformance traffic without any change to the DSCP service level.
- **Exceed** – Specifies whether traffic that exceeds the committed information 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 rate in kilobits per second. (Range: 0-10000000 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 in bytes.
(Range: 64-16000000 at a granularity of 4k bytes)
The burst size cannot exceed 16 Mbytes.
- **Peak Information Rate (PIR)** – Rate in kilobits per second.
(Range: 0-1000000 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)** – Burst size in bytes.
(Range: 64-16000000 at a granularity of 4k bytes)
The burst size cannot exceed 16 Mbytes.
- **Conform** – Specifies that traffic conforming to the committed information rate (CIR) and peak burst size (BP) will be transmitted without any change to the DSCP service level.
 - **Transmit** – Transmits in-conformance traffic without any change to the DSCP service level.
- **Exceed** – Specifies whether traffic that exceeds the committed information rate (CIR) or committed burst size (BC) 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 151: 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 152: Showing Policy Maps



Traffic > DiffServ

Step: 2. Configure Policy Action: Show

Policy List Total: 1

	Policy Name	Description
☐	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 153: Adding Rules to a Policy Map

Traffic > DiffServ

Step: 2. Configure Policy Action: Add Rule

Policy Name: rd-policy

Rule:

Class Name: rd-class

☒ Action: Set PHB (0-7) 3

☒ Meter

Meter Mode: Flow

Committed Information Rate (0-10000000): 1000000 kbps

Committed Burst Size (64-16000000): 4000 bytes

Excess Burst Size (64-16000000): bytes

Peak Information Rate (0-10000000): kbps

Peak Burst Size (64-16000000): bytes

Conform: Transmit

Exceed: Set IP DSCP (0-63)

Violate: Drop

Apply Revert

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 154: Showing the Rules for a Policy Map

The screenshot shows the 'Traffic > DiffServ' configuration page. The 'Step' is set to '2. Configure Policy' and the 'Action' is 'Show Rule'. The 'Policy Name' is 'rd-policy'. Below this, the 'Rule List' is displayed with a total of 1 rule. The table has columns for Class Name, Action, Meter Mode, Committed Information Rate (kbps), Committed Burst Size (bytes), Exceeded Burst Size (bytes), Peak Information Rate (kbps), Peak Burst Size (bytes), Conform, Exceed, and Violate. The rule shown is 'rd-class' with action 'Set PHB 3', meter mode 'Flow', committed rate '1000000', and committed burst size '4000'. The 'Conform' action is 'Transmit' and the 'Violate' action is 'Drop'.

Class Name	Action	Meter								
		Meter Mode	Committed Information Rate (kbps)	Committed Burst Size (bytes)	Exceeded Burst Size (bytes)	Peak Information Rate (kbps)	Peak Burst Size (bytes)	Conform	Exceed	Violate
rd-class	Set PHB 3	Flow	1000000	4000				Transmit		Drop

Buttons: Delete, Revert

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 then bind the service policy to the required interface.

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 155: Attaching a Policy Map to a Port

Traffic > DiffServ

Step: 3. Configure Interface

Port Service Policy List Total: 28

Port	Ingress	Egress
1	☐ rd-policy	☐ rd-policy
2	☐ rd-policy	☐ rd-policy
3	☐ rd-policy	☐ rd-policy
4	☐ rd-policy	☐ rd-policy
5	☐ rd-policy	☐ rd-policy
6	☐ rd-policy	☐ rd-policy
7	☐ rd-policy	☐ rd-policy
8	☐ rd-policy	☐ rd-policy
9	☐ rd-policy	☐ rd-policy
10	☐ rd-policy	☐ rd-policy

Apply Revert

VoIP Traffic Configuration

This chapter covers the following topics:

- ◆ **Global Settings** – Enables VOIP globally, sets the Voice VLAN, and the aging time for attached ports.
- ◆ **Telephony OUI List** – Configures the list of phones to be treated as VOIP devices based on the specified Organization Unit Identifier (OUI).
- ◆ **Port Settings** – Configures the way in which a port is added to the Voice VLAN, the filtering of non-VoIP packets, the method of detecting VoIP traffic, and the priority assigned to voice traffic.

Overview

When IP telephony is deployed in an enterprise network, it is recommended to isolate the Voice over IP (VoIP) network traffic from other data traffic. Traffic isolation can provide higher voice quality by preventing excessive packet delays, packet loss, and jitter. This is best achieved by assigning all VoIP traffic to a single Voice VLAN.

The use of a Voice VLAN has several advantages. It provides security by isolating the VoIP traffic from other data traffic. End-to-end QoS policies and high priority can be applied to VoIP VLAN traffic across the network, guaranteeing the bandwidth it needs. VLAN isolation also protects against disruptive broadcast and multicast traffic that can seriously affect voice quality.

The switch allows you to specify a Voice VLAN for the network and set a CoS priority for the VoIP traffic. The VoIP traffic can be detected on switch ports by using the source MAC address of packets, or by using LLDP (IEEE 802.1AB) to discover connected VoIP devices. When VoIP traffic is detected on a configured port, the switch automatically assigns the port as a tagged member the Voice VLAN. Alternatively, switch ports can be manually configured.

Configuring VoIP Traffic

Use the Traffic > VoIP (Configure Global) page to configure the switch for VoIP traffic. First enable automatic detection of VoIP devices attached to the switch ports, then set the Voice VLAN ID for the network. The Voice VLAN aging time can also be set to remove a port from the Voice VLAN when VoIP traffic is no longer received on the port.

Command Usage

All ports are set to VLAN hybrid mode by default. Prior to enabling VoIP for a port (by setting the VoIP mode to Auto or Manual as described below), first ensure that VLAN membership is not set to access mode (see [“Adding Static Members to VLANs” on page 165](#)).

Parameters

These parameters are displayed:

- ◆ **Auto Detection Status** – Enables the automatic detection of VoIP traffic on switch ports. (Default: Disabled)
- ◆ **Voice VLAN** – Sets the Voice VLAN ID for the network. Only one Voice VLAN is supported and it must already be created on the switch. (Range: 1-4094)
- ◆ **Voice VLAN Aging Time** – The time after which a port is removed from the Voice VLAN when VoIP traffic is no longer received on the port. (Range: 5-43200 minutes; Default: 1440 minutes)



Note: The Voice VLAN ID cannot be modified when the global Auto Detection Status is enabled.

Web Interface

To configure global settings for a Voice VLAN:

1. Click Traffic, VoIP.
2. Select Configure Global from the Step list.
3. Enable Auto Detection.
4. Specify the Voice VLAN ID.
5. Adjust the Voice VLAN Aging Time if required.
6. Click Apply.

Figure 156: Configuring a Voice VLAN

The screenshot shows a web interface for configuring VoIP settings. At the top, it says 'Traffic > VoIP'. Below that is a 'Step:' dropdown menu currently showing '1. Configure Global'. The main configuration area has three rows: 'Auto Detection Status' with a checked checkbox and the text 'Enabled'; 'Voice VLAN' with a dropdown menu showing '1234'; and 'Voice VLAN Aging Time (5-43200)' with a text input field containing '3000' followed by 'sec'. At the bottom right of the form are two buttons: 'Apply' and 'Revert'.

Configuring Telephony OUI

VoIP devices attached to the switch can be identified by the vendor's Organizational Unique Identifier (OUI) in the source MAC address of received packets. OUI numbers are assigned to vendors and form the first three octets of device MAC addresses. The MAC OUI numbers for VoIP equipment can be configured on the switch so that traffic from these devices is recognized as VoIP. Use the Traffic > VoIP (Configure OUI) page to configure this feature.

Parameters

These parameters are displayed:

- ◆ **Telephony OUI** – Specifies a MAC address range to add to the list.
(Format: xx-xx-xx-xx-xx-xx)
- ◆ **Mask** – Identifies a range of MAC addresses. Setting a mask of FF-FF-FF-00-00-00 identifies all devices with the same OUI (the first three octets). Other masks restrict the MAC address range. Setting a mask of FF-FF-FF-FF-FF-FF specifies a single MAC address. (Format: xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx;
Default: FF-FF-FF-00-00-00)
- ◆ **Description** – User-defined text that identifies the VoIP devices.

Web Interface

To configure MAC OUI numbers for VoIP equipment:

1. Click Traffic, VoIP.
2. Select Configure OUI from the Step list.
3. Select Add from the Action list.
4. Enter a MAC address that specifies the OUI for VoIP devices in the network.
5. Select a mask from the pull-down list to define a MAC address range.

6. Enter a description for the devices.
7. Click Apply.

Figure 157: Configuring an OUI Telephony List

Traffic > VoIP

Step: 2. Configure OUI Action: Add

Telephony OUI: 00-e0-bb-00-00-00

Mask: FF-FF-FF-00-00-00

Description: old phones

Apply Revert

To show the MAC OUI numbers used for VoIP equipment:

1. Click Traffic, VoIP.
2. Select Configure OUI from the Step list.
3. Select Show from the Action list.

Figure 158: Showing an OUI Telephony List

Traffic > VoIP

Step: 2. Configure OUI Action: Show

Telephony OUI List Total: 3

☐	Telephony OUI	Mask	Description
☐	00-E0-BB-00-00-00	FF-FF-FF-00-00-00	old phones
☐	00-11-22-33-44-55	FF-FF-FF-00-00-00	new phones
☐	00-98-76-54-32-10	FF-FF-FF-FF-FF-FF	Chris' phone

Delete Revert

Configuring VoIP Traffic Ports

Use the Traffic > VoIP (Configure Interface) page to configure ports for VoIP traffic, you need to set the mode (Auto or Manual), specify the discovery method to use, and set the traffic priority. You can also enable security filtering to ensure that only VoIP traffic is forwarded on the Voice VLAN.

Command Usage

All ports are set to VLAN hybrid mode by default. Prior to enabling VoIP for a port (by setting the VoIP mode to Auto or Manual as described below), first ensure that VLAN membership is not set to access mode (see [“Adding Static Members to VLANs” on page 165](#)).

Parameters

These parameters are displayed:

- ◆ **Mode** – Specifies if the port will be added to the Voice VLAN when VoIP traffic is detected. (Default: None)
 - **None** – The Voice VLAN feature is disabled on the port. The port will not detect VoIP traffic or be added to the Voice VLAN.
 - **Auto** – The port will be added as a tagged member to the Voice VLAN when VoIP traffic is detected on the port. You must select a method for detecting VoIP traffic, either OUI or 802.1AB (LLDP). When OUI is selected, be sure to configure the MAC address ranges in the Telephony OUI list.
 - **Manual** – The Voice VLAN feature is enabled on the port, but the port must be manually added to the Voice VLAN.
- ◆ **Security** – Enables security filtering that discards any non-VoIP packets received on the port that are tagged with the voice VLAN ID. VoIP traffic is identified by source MAC addresses configured in the Telephony OUI list, or through LLDP that discovers VoIP devices attached to the switch. Packets received from non-VoIP sources are dropped. (Default: Disabled)
- ◆ **Discovery Protocol** – Selects a method to use for detecting VoIP traffic on the port. (Default: OUI)
 - **OUI** – Traffic from VoIP devices is detected by the Organizationally Unique Identifier (OUI) of the source MAC address. OUI numbers are assigned to vendors and form the first three octets of a device MAC address. MAC address OUI numbers must be configured in the Telephony OUI list so that the switch recognizes the traffic as being from a VoIP device.
 - **LLDP** – Uses LLDP (IEEE 802.1AB) to discover VoIP devices attached to the port. LLDP checks that the “telephone bit” in the system capability TLV is turned on. See [“Link Layer Discovery Protocol” on page 389](#) for more information on LLDP.
- ◆ **Priority** – Defines a CoS priority for port traffic on the Voice VLAN. The priority of any received VoIP packet is overwritten with the new priority when the Voice VLAN feature is active for the port. (Range: 0-6; Default: 6)
- ◆ **Remaining Age** – Number of minutes before this entry is aged out.

The Remaining Age starts to count down when the OUI’s MAC address expires from the MAC address table. Therefore, the MAC address aging time should be added to the overall aging time. For example, if you configure the MAC address table aging time to 30 seconds, and the voice VLAN aging time to 5 minutes, then after 5.5 minutes, a port will be removed from voice VLAN when VoIP traffic is no longer received on the port. Alternatively, if you clear the MAC address table manually, then the switch will also start counting down the Remaining Age.


When VoIP Mode is set to Auto, the Remaining Age will be displayed.
Otherwise, if the VoIP Mode is Disabled or set to Manual, the remaining age will display "NA."

Web Interface

To configure VoIP traffic settings for a port:

1. Click Traffic, VoIP.
2. Select Configure Interface from the Step list.
3. Configure any required changes to the VoIP settings each port.
4. Click Apply.

Figure 159: Configuring Port Settings for a Voice VLAN



Port	Mode	Security	Discovery Protocol	Priority (0-6)	Remaining Age (minutes)
1	None	☒ Enabled	☒ OUI ☐ LLDP	6	NA
2	None	☒ Enabled	☒ OUI ☐ LLDP	6	NA
3	Manual	☒ Enabled	☒ OUI ☐ LLDP	5	NA
4	None	☒ Enabled	☒ OUI ☐ LLDP	6	NA
5	None	☐ Enabled	☒ OUI ☐ LLDP	6	NA

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, specify authentication servers, configure remote authentication and accounting.
- ◆ [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.
- ◆ [DoS Protection](#) – Protects against Denial-of-Service attacks.
- ◆ [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:

- ◆ Accounting for IEEE 802.1X authenticated users that access the network through the switch.
- ◆ Accounting for users that access management interfaces on the switch through the console and Telnet.
- ◆ Accounting for commands that users enter at specific CLI privilege levels.
- ◆ Authorization 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 275](#).
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 160: 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 161: 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

- ◆ **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.)
 - **Accounting Server UDP Port** – Network (UDP) port on authentication server used for accounting messages. (Range: 1-65535; Default: 1813)
 - **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: 1-65535; Default: 5)
 - **Authentication Retries** – Number of times the switch tries to authenticate logon access via the authentication server. (Range: 1-30; 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: 1-65535; Default: 5)
- **Authentication Retries** – Number of times the switch tries to authenticate logon access via the authentication server. (Range: 1-30; 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.

Configure Group

- ◆ **Server Type** – Select RADIUS or TACACS+ server.
- ◆ **Group Name** – Defines a name for the RADIUS or TACACS+ server group. (Range: 1-64 characters)
- ◆ **Sequence at Priority** – Specifies the server and sequence to use for the group. (Range: 1-5 for RADIUS; 1 for TACACS)

When specifying the priority sequence for a sever, the server index must already be defined (see [“Configuring Local/Remote Logon Authentication” on page 275](#)).

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 162: Configuring Remote Authentication Server (RADIUS)

The screenshot shows the 'Security > AAA > Server' configuration page. The 'Server Type' is set to 'RADIUS'. Under 'Server Index', 'Global' is selected. The configuration fields are as follows:

Field	Value
Server IP Address	10.1.1.1
Accounting Server UDP Port (1-65535)	1813
Authentication Server UDP Port (1-65535)	1815
Authentication Timeout (1-65535)	10 sec
Authentication Retries (1-30)	5
Set Key	☒
Authentication Key	
Confirm Authentication Key	

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

Figure 163: Configuring Remote Authentication Server (TACACS+)

Security > AAA > Server

Step: 1. Configure Server

Server Type ☐ RADIUS ☒ TACACS+

☐ Global | Server Index: ☒ 1

Server IP Address

Authentication Server TCP Port (1-65535)

Authentication Timeout (1-540) sec

Authentication Retries (1-30)

☒ Set Key

Authentication Key

Confirm Authentication Key

Apply Revert

To configure the RADIUS or TACACS+ server groups to use for accounting and authorization:

1. Click Security, AAA, Server.
2. Select Configure Group from the Step list.
3. Select Add from the Action list.
4. Select RADIUS or TACACS+ server type.
5. Enter the group name, followed by the index of the server to use for each priority level.
6. Click Apply.

Figure 164: Configuring AAA Server Groups

Security > AAA > Server

Step: 2. Configure Group Action: Add

Server Type ☒ RADIUS ☐ TACACS+

RADIUS Group Name

Sequence At Priority 1

Sequence At Priority 2

Sequence At Priority 3

Sequence At Priority 4

Sequence At Priority 5

Apply Revert

To show the RADIUS or TACACS+ server groups used for accounting and authorization:

1. Click Security, AAA, Server.
2. Select Configure Group from the Step list.
3. Select Show from the Action list.

Figure 165: Showing AAA Server Groups

The screenshot shows the 'Security > AAA > Server' configuration page. At the top, there is a breadcrumb trail. Below it, a 'Step:' dropdown is set to '2. Configure Group' and an 'Action:' dropdown is set to 'Show'. Under 'Server Type', the 'RADIUS' radio button is selected. The 'RADIUS Group List' section shows a total of 3 groups. A table lists the groups with their names and member indices. At the bottom right, there are 'Delete' and 'Revert' buttons.

	Group Name	Member Index
☐	radius	1, 2, 3, 5
☐	radius1	3, 5, 1
☐	radius2	1, 2, 5

Configuring AAA Accounting

Use the Security > AAA > Accounting page to enable accounting of requested services for billing or security purposes, and also to display the configured accounting methods, the methods applied to specific interfaces, and basic accounting information recorded for user sessions.

Command Usage

AAA authentication through a RADIUS or TACACS+ server must be enabled before accounting is enabled.

Parameters

These parameters are displayed:

Configure Global

- ◆ **Periodic Update** - Specifies the interval at which the local accounting service updates information for all users on the system to the accounting server. (Range: 1-2147483647 minutes)

Configure Method

- ◆ **Accounting Type** – Specifies the service as:
 - **802.1X** – Accounting for end users.
 - **Command** – Administrative accounting to apply to commands entered at specific CLI privilege levels.

- **Exec** – Administrative accounting for local console, Telnet, or SSH connections.
- ◆ **Privilege Level** – The CLI privilege levels (0-15). This parameter only applies to Command accounting.
- ◆ **Method Name** – Specifies an accounting method for service requests. The “default” methods are used for a requested service if no other methods have been defined. (Range: 1-64 characters)

Note that the method name is only used to describe the accounting method configured on the specified RADIUS or TACACS+ servers. No information is sent to the servers about the method to use.
- ◆ **Accounting Notice** – Records user activity from log-in to log-off point.
- ◆ **Server Group Name** - Specifies the accounting server group. (Range: 1-64 characters)

The group names “radius” and “tacacs+” specifies all configured RADIUS and TACACS+ hosts (see [“Configuring Local/Remote Logon Authentication” on page 275](#)). Any other group name refers to a server group configured on the Security > AAA > Server (Configure Group) page.

Configure Service

- ◆ **Accounting Type** – Specifies the service as 802.1X, Command or Exec as described in the preceding section.
- ◆ **802.1X**
 - **Method Name** – Specifies a user defined accounting method to apply to an interface. This method must be defined in the Configure Method page. (Range: 1-64 characters)
- ◆ **Command**
 - **Privilege Level** – The CLI privilege levels (0-15).
 - **Console Method Name** – Specifies a user-defined method name to apply to commands entered at the specified CLI privilege level through the console interface.
 - **VTY Method Name** – Specifies a user-defined method name to apply to commands entered at the specified CLI privilege level through Telnet or SSH.
- ◆ **Exec**
 - **Console Method Name** – Specifies a user defined method name to apply to console connections.

- **VTY Method Name** – Specifies a user defined method name to apply to Telnet and SSH connections.

Show Information – Summary

- ◆ **Accounting Type** - Displays the accounting service.
- ◆ **Method Name** - Displays the user-defined or default accounting method.
- ◆ **Server Group Name** - Displays the accounting server group.
- ◆ **Interface** - Displays the port, console or Telnet interface to which these rules apply. (This field is null if the accounting method and associated server group has not been assigned to an interface.)

Show Information – Statistics

- ◆ **User Name** - Displays a registered user name.
- ◆ **Accounting Type** - Displays the accounting service.
- ◆ **Interface** - Displays the receive port number through which this user accessed the switch.
- ◆ **Time Elapsed** - Displays the length of time this entry has been active.

Web Interface

To configure global settings for AAA accounting:

1. Click Security, AAA, Accounting.
2. Select Configure Global from the Step list.
3. Enter the required update interval.
4. Click Apply.

Figure 166: Configuring Global Settings for AAA Accounting

The screenshot shows a web interface for configuring AAA Accounting. The breadcrumb trail at the top is "Security > AAA > Accounting". Below this, there is a "Step:" dropdown menu currently set to "1. Configure Global". The main configuration area contains a label "Periodic Update (1-2147483647)" followed by a checked checkbox and a text input field containing the number "1", with the unit "min" to its right. At the bottom right of the form are two buttons: "Apply" and "Revert".

To configure the accounting method applied to various service types and the assigned server group:

1. Click Security, AAA, Accounting.
2. Select Configure Method from the Step list.
3. Select Add from the Action list.
4. Select the accounting type (802.1X, Command, Exec).
5. Specify the name of the accounting method and server group name.
6. Click Apply.

Figure 167: Configuring AAA Accounting Methods

The screenshot displays the 'Security > AAA > Accounting' configuration page. At the top, the breadcrumb trail is visible. Below it, the 'Step' dropdown is set to '2. Configure Method' and the 'Action' dropdown is set to 'Add'. The main configuration area contains the following fields:

- Accounting Type:** A dropdown menu with '802.1X' selected.
- Method Name:** A text input field containing 'default'.
- Accounting Notice:** A dropdown menu with 'Start-Stop' selected.
- Server Group Name:** A section with two radio buttons. The first radio button is selected and is followed by a dropdown menu showing 'radius'. The second radio button is unselected and is followed by an empty text input field.

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

To show the accounting method applied to various service types and the assigned server group:

1. Click Security, AAA, Accounting.
2. Select Configure Method from the Step list.
3. Select Show from the Action list.

Figure 168: Showing AAA Accounting Methods

Security > AAA > Accounting

Step: 2. Configure Method Action: Show

Method List Total: 18

Accounting Type	Method Name	Accounting Notice	Server Group Name
802.1X	default	Start-Stop	radius
Command 0	default	Start-Stop	tacacs+
Command 1	default	Start-Stop	tacacs+
Command 2	default	Start-Stop	tacacs+
Command 3	default	Start-Stop	tacacs+
Command 4	default	Start-Stop	tacacs+
Command 5	default	Start-Stop	tacacs+
Command 6	default	Start-Stop	tacacs+
Command 7	default	Start-Stop	tacacs+
Command 8	default	Start-Stop	tacacs+

Delete Revert

To configure the accounting method applied to specific interfaces, console commands entered at specific privilege levels, and local console, Telnet, or SSH connections:

1. Click Security, AAA, Accounting.
2. Select Configure Service from the Step list.
3. Select the accounting type (802.1X, Command, Exec).
4. Enter the required accounting method.
5. Click Apply.

Figure 169: Configuring AAA Accounting Service for 802.1X Service

Security > AAA > Accounting

Step: 3. Configure Service

Accounting Type ☒ 802.1X ☐ Command ☐ EXEC

Port Method List Total: 28

Port	Method Name
1	default
2	default
3	default
4	
5	

Figure 170: Configuring AAA Accounting Service for Command Service

Security > AAA > Accounting

Step: 3. Configure Service

Accounting Type: ☐ 802.1X ☒ Command ☐ EXEC

Command Method List Total: 16

Privilege Level	Console Method Name	VTY Method Name
0	default	default
1	default	default
2	default	default
3	default	default
4	command4Method	default
5	default	command5Method

Figure 171: Configuring AAA Accounting Service for Exec Service

Security > AAA > Accounting

Step: 3. Configure Service

Accounting Type: ☐ 802.1X ☐ Command ☒ EXEC

Console Method Name: default

VTY Method Name: default

Apply Revert

To display a summary of the configured accounting methods and assigned server groups for specified service types:

1. Click Security, AAA, Accounting.
2. Select Show Information from the Step list.
3. Click Summary.

Figure 172: Displaying a Summary of Applied AAA Accounting Methods

Security > AAA > Accounting

Step: 4. Show Information

☒ Summary ☐ Statistics

Method List Total: 18

Accounting Type	Method Name	Server Group Name	Interface
802.1X	default	radius	
Command 0	default	tacacs+	
Command 1	default	tacacs+	
Command 2	default	tacacs+	
Command 3	default	tacacs+	
Command 4	default	tacacs+	
Command 5	default	tacacs+	
Command 6	default	tacacs+	
Command 7	default	tacacs+	
Command 8	default	tacacs+	

To display basic accounting information and statistics recorded for user sessions:

1. Click Security, AAA, Accounting.
2. Select Show Information from the Step list.
3. Click Statistics.

Figure 173: Displaying Statistics for AAA Accounting Sessions

Security > AAA > Accounting

Step: 4. Show Information

☐ Summary ☒ Statistics

Accounting Statistics Total: 2

User Name	Accounting Type	Interface	Time Elapsed
Bob	802.1X	Eth1/1	3:44:55
Ted	802.1X	Eth1/5	1:24:51

Configuring AAA Authorization

Use the Security > AAA > Authorization page to enable authorization of requested services, and also to display the configured authorization methods, and the methods applied to specific interfaces.

Command Usage

- ◆ This feature performs authorization to determine if a user is allowed to run an Exec shell.
- ◆ AAA authentication through a RADIUS or TACACS+ server must be enabled before authorization is enabled.

Parameters

These parameters are displayed:

Configure Method

- ◆ **Authorization Type** – Specifies the service as Exec, indicating administrative authorization for local console, Telnet, or SSH connections.
- ◆ **Method Name** – Specifies an authorization method for service requests. The “default” method is used for a requested service if no other methods have been defined. (Range: 1-64 characters)
- ◆ **Server Group Name** - Specifies the authorization server group. (Range: 1-64 characters)

The group name “tacacs+” specifies all configured TACACS+ hosts (see [“Configuring Local/Remote Logon Authentication” on page 275](#)). Any other group name refers to a server group configured on the TACACS+ Group Settings page. Authorization is only supported for TACACS+ servers.

Configure Service

- ◆ **Authorization Type** – Specifies the service as Exec, indicating administrative authorization for local console, Telnet, or SSH connections.
- ◆ **Console Method Name** – Specifies a user defined method name to apply to console connections.
- ◆ **VTY Method Name** – Specifies a user defined method name to apply to Telnet and SSH connections.

Show Information

- ◆ **Authorization Type** - Displays the authorization service.
- ◆ **Method Name** - Displays the user-defined or default accounting method.
- ◆ **Server Group Name** - Displays the authorization server group.
- ◆ **Interface** - Displays the console or Telnet interface to which these rules apply. (This field is null if the authorization method and associated server group has not been assigned to an interface.)

Web Interface

To configure the authorization method applied to the Exec service type and the assigned server group:

1. Click Security, AAA, Authorization.
2. Select Configure Method from the Step list.

3. Specify the name of the authorization method and server group name.
4. Click Apply.

Figure 174: Configuring AAA Authorization Methods

Security > AAA > Authorization

Step: 1. Configure Method Action: Add

Authorization Type: EXEC

Method Name: default

Server Group Name: ☒ tacacs+ ☐

Apply Revert

To show the authorization method applied to the EXEC service type and the assigned server group:

1. Click Security, AAA, Authorization.
2. Select Configure Method from the Step list.
3. Select Show from the Action list.

Figure 175: Showing AAA Authorization Methods

Security > AAA > Authorization

Step: 1. Configure Method Action: Show

Method List Total: 2

	Authorization Type	Method Name	Server Group Name
☐	EXEC	default	tacacs+
☐	EXEC	aaa	tacacs1

Delete Revert

To configure the authorization method applied to local console, Telnet, or SSH connections:

1. Click Security, AAA, Authorization.
2. Select Configure Service from the Step list.
3. Enter the required authorization method.
4. Click Apply.

Figure 176: Configuring AAA Authorization Methods for Exec Service

Security > AAA > Authorization

Step: 2. Configure Service

Authorization Type: EXEC

Console Method Name: tps-auth

VTY Method Name: tps-auth

Apply Revert

To display a the configured authorization method and assigned server groups for The Exec service type:

1. Click Security, AAA, Authorization.
2. Select Show Information from the Step list.

Figure 177: Displaying the Applied AAA Authorization Method

Security > AAA > Authorization

Step: 3. Show Information

Method List Total: 3

Authorization Type	Method Name	Server Group Name	Interface
EXEC	default	tacacs+	
EXEC	console	tacacs+	Console
EXEC	telnet	tacacs+	Telnet

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 command access privileges. (Range: 0-15)

Level 0, 8 and 15 are designed for users (guest), managers (network maintenance), and administrators (top-level access). The other levels can be used to configured specialized access profiles.

Level 0-7 provide the same default access to a limited number of commands which display the current status of the switch, as well as several database clear and reset functions. These commands are equivalent to those available under Normal Exec command mode in the CLI.

Level 8-14 provide the same default access privileges, including additional commands beyond those provided for Levels 0-7 (equivalent to CLI Normal Exec command mode), and a subset of the configuration commands provided for Level 15 (equivalent to CLI Privileged Exec command mode).

Level 15 provides full access to all commands.

The privilege level associated with any command can be changed using the “privilege” command described in the *CLI Reference Guide*.

Any privilege level can access all of the commands assigned to lower privilege levels. For example, privilege level 8 can access all commands assigned to privilege levels 7-0 according to default settings, and to any other commands assigned to levels 7-0 using the “privilege” command described in the *CLI Reference Guide*.

◆ **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 178: 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 179: Showing User Accounts

Security > User Accounts

Action: Show

User Account List Total: 3

	User Name	Access Level
☐	admin	15
☐	guest	0
☐	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 275.](#))

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 180: Configuring Global Settings for Web Authentication

Security > Web Authentication

Step: 1. Configure Global

Web Authentication Status ☐ Enabled

Session Timeout (300-3600) 3600 sec

Quiet Period (1-180) 60 sec

Login Attempts (1-3) 3

Apply Revert

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 181: Configuring Interface Settings for Web Authentication

Security > Web Authentication

Steps: 2. Configure Interface

Port: 1

Status: ☒ Enabled

Apply Revert

Authenticated Host List Total: 2

	Host IP Address	Remaining Session Time (sec)
☐	10.1.1.1	300
☐	10.2.2.2	100

Re-authenticate Revert

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 276.)

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 18: 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=ip4acl
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 351](#)).

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 182: Configuring Global Settings for Network Access

Security > Network Access

Step: 1. Configure Global

Aging Status ☒ Enabled

Reauthentication Time (120-1000000) sec

Apply Revert

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 Count⁶** – 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 Count⁶** – 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: 1-2048; 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 162](#)). 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 351](#)).

A port can only be assigned to the guest VLAN if it failed authentication, and switchport mode is set to Hybrid. (See [“Adding Static Members to VLANs” on page 165](#).)
- ◆ **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)

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.

- ◆ **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)

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 183: Configuring Interface Settings for Network Access

Security > Network Access

Step: 2. Configure Interface

General Link Detection

Port List Total: 28

Port	MAC Authentication			Network Access Max MAC Count (1-2048)	Guest VLAN (0-4093, 0: Disabled)	Dynamic VLAN	Dynamic QoS	MAC Filter ID (1-64)
	Status	Intrusion	Max MAC Count (1-1024)					
1	☐ Enabled	Block	1024	1024	0	☒ Enabled	☐ Enabled	<input type="text"/>
2	☒ Enabled	Block	1024	1024	0	☒ Enabled	☐ Enabled	<input type="text"/>
3	☐ Enabled	Block	1024	1024	0	☒ Enabled	☐ Enabled	<input type="text"/>
4	☐ Enabled	Block	1024	1024	0	☒ Enabled	☐ Enabled	<input type="text"/>
5	☐ Enabled	Block	1024	1024	0	☒ Enabled	☐ Enabled	<input type="text"/>

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 184: Configuring Link Detection for Network Access

Security > Network Access

Step: 2. Configure Interface

General Link Detection

Port List Total: 28

Port	Link Detection Status	Condition	Action
1	☐ Enabled	Link down	Trap
2	☒ Enabled	Link up and down	Trap
3	☐ Enabled	Link down	Trap
4	☐ Enabled	Link down	Trap
5	☐ 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 185: Configuring a MAC Address Filter for Network Access

Security > Network Access

Step: 3. Configure MAC Filter Action: Add

Filter ID (1-64) 22

MAC Address 11-22-33-44-55-66


MAC Address Mask FFFFFFFF

Apply 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 186: 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 section is titled 'MAC Filter List' with a 'Total: 1' indicator. It contains a table with the following data:

☐	Filter ID	MAC Address	MAC Address Mask
☐	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 187: Showing Addresses Authenticated for Network Access

Security > Network Access

Step: 4. Show Information

Query by:

Sort Key: MAC Address

☐ MAC Address

☐ Interface

☐ Attribute

Query

Authenticated MAC Address List Total: 8

	MAC Address	Interface	RADIUS Server	Time	Attribute
☐	00-00-86-45-F2-23	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 16m 12s	Dynamic
☐	00-00-E8-5E-E1-DD	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 32m 24s	Dynamic
☐	00-00-E8-B1-93-30	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 40m 32s	Dynamic
☐	00-01-80-31-B8-30	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 18m 51s	Dynamic
☐	00-01-80-36-95-D8	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 32m 22s	Dynamic
☐	00-01-80-3B-D3-7F	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 22m 28s	Dynamic
☐	00-01-80-3C-3C-19	Unit 2 / Port 23	10.2.2.10	2008y 20m 12d 11h 15m 19s	Dynamic
☐	00-01-80-3C-3E-B3	Unit 2 / Port 23	10.2.2.10	2008y 20m 12d 11h 17m 40s	Dynamic

Delete Revert

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 6, Mozilla Firefox 4, or Google Chrome 29, or more recent versions.
- ◆ The following web browsers and operating systems currently support HTTPS:

Table 19: HTTPS System Support

Web Browser	Operating System
Internet Explorer 8.x or later	Windows 98, Windows NT (with service pack 6a), Windows 2000, XP, Vista, 7, 8
Mozilla Firefox 4 or later	Windows 2000, XP, Vista, 7, 8, Linux
Google Chrome 40 or later	Windows XP, Vista, 7, 8

- ◆ To specify a secure-site certificate, see [“Replacing the Default Secure-site Certificate” on page 306](#).



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. (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 188: Configuring HTTPS



Security > HTTPS

Action: Configure Global

HTTPS Status ☒ Enabled

UDP Port (1-65535)

Apply Revert

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 111](#) 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 189: Downloading the Secure-Site Certificate

The screenshot shows a web-based configuration interface for HTTPS settings. At the top, it says "Security > HTTPS". Below this, there is a section titled "Action:" with a dropdown menu set to "Copy Certificate". Underneath, there are five input fields: "TFTP Server IP Address" with the value "192.168.0.4", "Certificate Source File Name" with "ES3510MA-site-certificate", "Private Key Source File Name" with "ES3510MA-private-key", "Private Password" with masked characters "*****", and "Confirm Password" also with masked characters "*****". At the bottom right of the form are two buttons: "Apply" and "Revert".

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 275](#)). 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 313](#) 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 290](#).) 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 312](#).

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 190: Configuring the SSH Server

The screenshot shows a web interface for configuring the SSH server. The breadcrumb navigation at the top reads "Security > SSH". Below this, a "Step:" dropdown menu is set to "1. Configure Global". The configuration area contains the following fields:

SSH Server Status	☒ 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"/>

At the bottom right of the configuration area are two buttons: "Apply" and "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 313](#).



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 310](#).

Parameters

These parameters are displayed:

- ◆ **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.

- ◆ **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. Note that you must select this item from the Show page. (Default: Disabled)

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 191: Generating the SSH Host Key Pair

The screenshot shows the 'Security > SSH' configuration page. At the top, the breadcrumb 'Security > SSH' is displayed. Below it, the 'Step' dropdown is set to '2. Configure Host Key' and the 'Action' dropdown is set to 'Generate'. The 'Host-Key Type' dropdown is set to 'Both'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 192: 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 290](#)).
- ◆ **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 193: Copying the SSH User's Public Key

The screenshot shows a web interface for configuring SSH settings. The breadcrumb is "Security > SSH". The "Step:" dropdown is set to "3. Configure User Key" and the "Action:" dropdown is set to "Copy". The form contains the following fields:

User Name	steve
User-Key Type	RSA
TFTP Server IP Address	192.168.0.61
Source File Name	rsa.pub

At the bottom right of the form 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 194: Showing the SSH User's Public Key

Security > SSH

Step: 3. Configure User Key Action: Show

User Name: admin

Public-Key of User-Key

☐ RSA

```
1024 65537
1480220937772193109967882191897869076673078824151724992407025121828690162016474445339
1667942450570093339422932545152202043583189918588264701199521464739810320305865029839
8473670516025543210438660758858919850241166441730463579799255182511108846314033955899
732895735202234665421721768844978831196756189309582533
```

☐ DSA

```
ssh-dss
AAAAB3NzaC1kc3MAAACBAKM4i8EgUe89W+vh1+y4z12YpyK8cpCNz30rhCriv1ClKmdgE/fiZPsqYH/ClAB/
eJ1rNAX0sJTUxtb8e2GLk+x8UmQJVuSDdklwZ9ZRoFwWA10Yyi57V1TK3tUnT9ocCbVJNGekJkXdiuac4OP09
tAPEAuXBuAWGpDAGSmg6pHAAAFQCPzwn0rSaLiTq53jxtsj0RILZwAAABiL0Cr7GC7ARztGE9dRkT9oh9
uhuiAzcXrAcZTmxhjGsEtMlAtxKm+r1O6pnz2aX9KEzMewJEMuDphOTRnSpMv39XtSW1aSXWtKoGAcQgDsFqK
```

Clear

Access Control Lists

Access Control Lists (ACL) provide packet filtering for IPv4/IPv6 frames (based on address, protocol, Layer 4 protocol port number or TCP control code), IPv6 frames (based on address, DSCP traffic class, or next header type), 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 or egress 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 512.
- ◆ The maximum number of rules per system is 2048 rules.
- ◆ An ACL can have up to 2048 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 or egress 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.

Setting a Time Range Use the Security > ACL (Configure Time Range) page to set a time range during which ACL functions are applied.

Command Usage

If both an absolute rule and one or more periodic rules are configured for the same time range (i.e., named entry), that entry will only take effect if the current time is within the absolute time range and one of the periodic time ranges.

Parameters

These parameters are displayed:

Add

- ◆ **Time-Range Name** – Name of a time range. (Range: 1-16 characters)

Add Rule


- ◆ **Time-Range** – Name of a time range.
- ◆ **Mode**
 - **Absolute** – Specifies a specific time or time range.
 - **Start/End** – Specifies the hours, minutes, month, day, and year at which to start or end.
 - **Periodic** – Specifies a periodic interval.
 - **Start/To** – Specifies the days of the week, hours, and minutes at which to start or end.

Web Interface

To configure a time range:

1. Click Security, ACL.
2. Select Configure Time Range from the Step list.
3. Select Add from the Action list.
4. Enter the name of a time range.
5. Click Apply.

Figure 195: Setting the Name of a Time Range



The screenshot shows a web interface for configuring ACLs. At the top, it says "Security > ACL". Below that, there are two dropdown menus: "Step:" with "1. Configure Time-Range" selected, and "Action:" with "Add" selected. Below these, there is a text input field labeled "Time-Range Name" containing the text "R&D". At the bottom right, there are two buttons: "Apply" and "Revert".

To show a list of time ranges:

1. Click Security, ACL.
2. Select Configure Time Range from the Step list.

3. Select Show from the Action list.

Figure 196: Showing a List of Time Ranges

The screenshot shows the 'Security > ACL' configuration page. At the top, there is a breadcrumb 'Security > ACL'. Below it, a 'Step:' dropdown is set to '1. Configure Time-Range' and an 'Action:' dropdown is set to 'Show'. A section titled 'Time-Range List' shows 'Total: 1'. Below this is a table with two columns: a checkbox and 'Time-Range Name'. The first row has a checked checkbox and the name 'R&D'. At the bottom right, there are 'Delete' and 'Revert' buttons.

	Time-Range Name
☒	R&D

To configure a rule for a time range:

1. Click Security, ACL.
2. Select Configure Time Range from the Step list.
3. Select Add Rule from the Action list.
4. Select the name of time range from the drop-down list.
5. Select a mode option of Absolute or Periodic.
6. Fill in the required parameters for the selected mode.
7. Click Apply.

Figure 197: Add a Rule to a Time Range

The screenshot shows the 'Security > ACL' configuration page with the 'Step:' dropdown set to '1. Configure Time-Range' and the 'Action:' dropdown set to 'Add Rule'. The 'Time-Range' dropdown is set to 'R&D' and the 'Mode' dropdown is set to 'Periodic'. Below these, there are two columns for configuration: 'Start' and 'To'. The 'Start' column has 'Days of the week' set to 'Weekend', 'Hours (0-23)' set to '5', and 'Minutes (0-59)' set to '30'. The 'To' column has 'Days of the week' set to 'Sunday', and empty fields for 'Hours (0-23)' and 'Minutes (0-59)'. At the bottom right, there are 'Apply' and 'Revert' buttons.

Start		To	
Days of the week	Weekend	Days of the week	Sunday
Hours (0-23)	5	Hours (0-23)	
Minutes (0-59)	30	Minutes (0-59)	

To show the rules configured for a time range:

1. Click Security, ACL.
2. Select Configure Time Range from the Step list.

3. Select Show Rule from the Action list.

Figure 198: Showing the Rules Configured for a Time Range

The screenshot shows the 'Security > ACL' configuration page. At the top, there's a 'Step:' dropdown set to '1. Configure Time-Range' and an 'Action:' dropdown set to 'Show Rule'. Below this is a 'Time-Range' dropdown set to 'time1'. The main section is titled 'Time-Range Rule List' with a 'Total: 4' indicator. It contains a table with four columns: a checkbox, 'Mode', 'Start', and 'End'. The table lists four rules: an Absolute rule from 2009-01-01 10:05 to 2010-01-31 20:10, and three Periodic rules: Daily 10:05 to 20:10, Monday 10:05 to Tuesday 20:10, and Monday 00:00 to Tuesday 23:59. At the bottom right of the table are 'Delete' and 'Revert' buttons.

☐	Mode	Start	End
☐	Absolute	2009-01-01 10:05	2010-01-31 20:10
☐	Periodic	Daily 10:05	Daily 20:10
☐	Periodic	Monday 10:05	Tuesday 20:10
☐	Periodic	Monday 00:00	Tuesday 23:59

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 199: Showing TCAM Utilization



The screenshot shows the 'Security > ACL' configuration page. At the top, there are two dropdown menus: 'Step:' set to '2. Configure ACL' and 'Action:' set to 'Show TCAM'. Below these, a table displays the following statistics:

Total Policy Control Entries	6144
Free Policy Control Entries	6128
Entries Used by System	16
Entries Used by User	0
TCAM Utilization	0.26%

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: 16 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, as well as DSCP, and the next header type.
 - **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 336](#)).

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 200: Creating an ACL

Security > ACL

Step: 2. Configure ACL Action: Add

ACL Name: R&D

Type: IP Standard

Apply 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 201: Showing a List of ACLs

Security > ACL

Step: 2. Configure ACL Action: Show

ACL List Total: 10

	ACL Name	Type
☐	acIStandard1	IP Standard
☐	acIStandard2	IP Standard
☐	acIExtended1	IP Extended
☐	acIExtended2	IP Extended
☐	acIMAC1	MAC
☐	acIMAC2	MAC
☐	acIMAC2	MAC
☐	acIPv6Standard	IPv6 Standard
☐	acIPv6Extended	IPv6 Extended
☐	acIARP	ARP

Delete Revert

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.
- ◆ **Time Range** – Name of a time range.

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 202: Configuring a Standard IPv4 ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, it says 'Step: 2. Configure ACL' and 'Action: Add Rule'. Below this, there are radio buttons for 'Type': 'IP Standard' (selected), 'IP Extended', 'MAC', 'IPv6 Standard', 'IPv6 Extended', and 'ARP'. The 'Name' field is 'R&D'. The 'Action' dropdown is 'Permit'. The 'Address Type' dropdown is 'Host'. The 'Source IP Address' field is '10.1.1.21'. The 'Source Subnet Mask' field is '255.255.255.255'. There is a checked checkbox for 'Time-Range' and a dropdown for 'Time-Range' set to 'R&D'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 322](#).)
- ◆ **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:
 - **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

- ◆ **Time Range** – Name of a time range.

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 203: Configuring an Extended IPv4 ACL

The screenshot shows the 'Security > ACL' configuration page. The 'Step' is '2. 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'. The 'Source IP Address' is '10.7.1.0' and the 'Source Subnet Mask' is '255.255.255.0'. The 'Destination Address Type' is 'Any'. The 'Destination IP Address' is '0.0.0.0' and the 'Destination Subnet Mask' is '0.0.0.0'. The 'Protocol' is 'TCP (6)'. The 'Control Code (0-63)' is empty. The 'Control Code Bit Mask (0-63)' is empty. The 'Service Type' is 'Precedence (0-7)'. The 'Time-Range' is 'R&D'. The 'Apply' and 'Revert' buttons are at the bottom right.

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)
- ◆ **Time Range** – Name of a time range.

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 204: Configuring a Standard IPv6 ACL

The screenshot shows the 'Security > ACL' configuration page. The 'Step' is set to '2. Configure ACL' and the 'Action' is 'Add Rule'. Under 'Type', 'IPv6 Standard' is selected. The 'Name' is 'R&D#6S'. The 'Action' is 'Permit'. The 'Source Address Type' is 'Host'. The 'Source IPv6 Address' is '2009:DB9:2229::79'. The 'Source Prefix Length (0-128)' is '128'. The 'Time-Range' checkbox is unchecked. The 'R&D' dropdown is set to 'R&D'. The 'Apply' and 'Revert' buttons are at the bottom right.

Security > ACL	
Step:	2. Configure ACL
Action:	Add Rule
Type	☐ IP Standard ☐ IP Extended ☐ MAC ☒ IPv6 Standard ☐ IPv6 Extended ☐ ARP
Name	R&D#6S
Action	Permit
Source Address Type	Host
Source IPv6 Address	2009:DB9:2229::79
Source Prefix Length (0-128)	128
☐ Time-Range	R&D
<div>Apply Revert</div>	

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.
- ◆ **Source Address Type** – Specifies the source 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 Address Type** – Specifies the destination IP address type. Use “Any” to include all possible addresses, or “IPv6-Prefix” to specify a range of addresses. (Options: Any, IPv6-Prefix; Default: Any)
- ◆ **Source/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.
- ◆ **Source/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 for the source prefix; 0-8 bits for the destination prefix)
- ◆ **DSCP** – DSCP traffic class. (Range: 0-63)
- ◆ **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)

- ◆ **Time Range** – Name of a time range.

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 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. Set any other required criteria, such as DSCP or next header type.
10. Click Apply.

Figure 205: Configuring an Extended IPv6 ACL

The screenshot shows the 'Security > ACL' configuration page. The 'Step' is set to '2. Configure ACL' and the 'Action' is 'Add Rule'. Under the 'Type' section, 'IPv6 Extended' is selected. The 'Name' is 'ipv6e'. The 'Action' is 'Permit'. The 'Source Address Type' is 'IPv6-Prefix'. The 'Source IPv6 Address' is '2009:DB9:2229:79' and the 'Source Prefix Length (0-128)' is '8'. The 'Destination Address Type' is 'Any'. The 'Destination IPv6 Address' is '::' and the 'Destination Prefix Length (0-8)' is '0'. The 'DSCP (0-63)' and 'Next-Header (0-255)' fields are empty. The 'Time-Range' checkbox is unchecked, and the 'R&D' dropdown is set to 'R&D'. The 'Apply' and 'Revert' buttons are at the bottom right.

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)
- ◆ **Internet Protocol** – Layer 3 or 4 information to match.
 - **No** – Not applied.
 - **IPv4** – See [“Configuring an Extended IPv4 ACL” on page 323.](#)
 - **IPv6** – See [“Configuring an Extended IPv6 ACL” on page 327.](#)
- ◆ **Time Range** – Name of a time range.

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 206: Configuring a MAC ACL

The screenshot displays the 'Security > ACL' configuration page. At the top, 'Step: 2. Configure ACL' and 'Action: Add Rule' are selected. Under the 'Type' section, 'MAC' is selected with radio buttons, and the 'Name' is set to 'mac'. The 'Action' is set to 'Permit'. For 'Source Address Type', 'Any' is selected, with corresponding fields for 'Source MAC Address' (00-00-00-00-00-00) and 'Source Bit Mask' (00-00-00-00-00-00). Similarly, for 'Destination Address Type', 'Any' is selected, with fields for 'Destination MAC Address' (00-00-00-00-00-00) and 'Destination Bit Mask' (00-00-00-00-00-00). The 'Packet Format' is set to 'Any'. Fields for 'VID (1-4094)' and 'VID Bit Mask (0-4095)' are present but empty. Under 'Internet Protocol', 'No' is selected with radio buttons. There are checkboxes for 'Time-Range' (unchecked) and 'R&D' (checked with a dropdown arrow). At the bottom right, there are 'Apply' and 'Revert' buttons for both the main configuration and the 'Time-Range' section.

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 337](#)).

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 322](#).)
- ◆ **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 207: Configuring a ARP ACL

The screenshot shows the 'Security > ACL' configuration page. 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 'IP'. 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'. The 'Log' checkbox is unchecked. The 'Apply' and 'Revert' buttons are at the bottom right.

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.

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.
- ◆ **Time Range** – Name of a time range.

- ◆ **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 options.
5. Select a port.
6. Select the name of an ACL from the ACL list.
7. Click Apply.

Figure 208: Binding a Port to an ACL

The screenshot shows the 'Security > ACL' web interface. At the top, the 'Step' is set to '3. Configure Interface'. Below this, there are three radio buttons for 'Type': IP (selected), MAC, and IPv6. A 'Port' dropdown menu is set to '1'. Under the 'Ingress' section, there are three options: 'ACL' (checked with a dropdown showing 'ips'), 'Time-Range' (unchecked with a dropdown showing 'R&D'), and 'Counter' (unchecked). Under the 'Egress' section, there are three options: 'ACL' (unchecked with a dropdown showing 'ips'), 'Time-Range' (unchecked with a dropdown showing 'R&D'), and 'Counter' (unchecked). At the bottom right, there are 'Apply' and 'Revert' buttons.

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 332](#).
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 209: Configuring ACL Mirroring

The screenshot shows a web interface titled "Security > ACL". At the top, there are two dropdown menus: "Step:" with "3. Configure Interface" selected, and "Action:" with "Add Mirror" selected. Below these, there are two more dropdown menus: "Port:" with "1" selected, and "ACL:" with "ips" selected. At the bottom right of the form, there are two buttons: "Apply" and "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 210: Showing the VLANs to Mirror



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. (Range: 1-28/52)
- ◆ **Type** – Selects the type of ACL.
- ◆ **Direction** – Displays statistics for ingress or egress traffic.
- ◆ **Query** – Displays statistics for selected criteria.
- ◆ **ACL Name** – The ACL bound 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.

⁷. Due to a hardware limitation, statistics are only displayed for permit rules.

4. Select a port.
5. Select ingress or egress traffic.

Figure 211: Showing ACL Statistics

The screenshot shows the 'Security > ACL' configuration page. At the top, there's a 'Step: 3. Configure Interface' and an 'Action: Show Hardware Counter' dropdown. Below this, there are three dropdown menus: 'Port' set to '1', 'Type' set to 'IP Standard', and 'Direction' set to 'Ingress'. A 'Query' button is located to the right of these settings. Below the settings, the 'ACL Name' is 'rd'. A 'Total: 1' label is present above a table. The table has five columns: 'Action', 'Source IP Address', 'Time-Range', 'Hit', and 'Clear Counter'. The first row shows 'Permit' for the action, 'Any' for the source IP address, and '14' for the hit count. A 'Clear' button is located next to the hit count.

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 377](#)). 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 331](#)).

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 212: Configuring Global Settings for ARP Inspection

The screenshot shows the 'Security > ARP Inspection' configuration page. The 'Step' dropdown is set to '1. Configure General'. The configuration options are as follows:

ARP Inspection Status	☒ Enabled
ARP Inspection Validate	☒ Dst-MAC ☐ IP ☐ Src-MAC
Log Message Number (0-256)	<input type="text" value="50"/>
Log Interval (0-86400)	<input type="text" value="100"/> sec

At the bottom right, there are two buttons: 'Apply' and '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 331](#)).
- ◆ 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 213: Configuring VLAN Settings for ARP Inspection

Security > ARP Inspection

Step: 2. Configure VLAN

ARP Inspection VLAN ID: 1

ARP Inspection VLAN Status: ☒ Enabled

ARP Inspection ACL Name: ☒ aaa ☐ Static

Apply Revert

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 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 214: Configuring Interface Settings for ARP Inspection

Port	Trust Status	Packet Rate Limit (0-2048 pps)
1	☒ Enabled	☒ 15
2	☐ Enabled	☒ 15
3	☐ Enabled	☒ 15
4	☐ Enabled	☒ 15
5	☐ Enabled	☒ 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 20: 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.
ARP packets dropped by additional validation (IP)	Count of ARP packets that failed the IP address test.
ARP packets dropped by additional validation (Dst-MAC)	Count of packets that failed the destination MAC address test.
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 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 215: Displaying Statistics for ARP Inspection

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 21: 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 216: 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 217: Creating an IP Address Filter for Management Access

The screenshot displays the 'Security > IP Filter' configuration page. At the top, the breadcrumb 'Security > IP Filter' is shown. Below it, the 'Action:' dropdown menu is set to 'Add'. The 'Mode' section contains four radio buttons: 'Web' (selected), 'SNMP', 'Telnet', and 'All'. Below the mode selection, there are two text input fields: 'Start IP Address' with the value '10.1.2.3' and 'End IP Address' which is currently empty. At the bottom right of the form, there are two buttons: 'Apply' and '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 218: Showing IP Addresses Authorized for Management Access



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 195](#)).

- ◆ 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 116](#)).
- ◆ 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. 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 301](#).
- ◆ **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 219: 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	☒ Enabled	Secure/Down	Trap and Shutdown	0	1	Disabled	0	NA	NA
2	☒ Enabled	Secure/Down	Trap	0	0	Disabled	1	NA	NA
3	☒ 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 220: Configuring Port Authentication



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 8, 7, Vista and XP, and in Windows 2000 with Service Pack 4. To support these encryption methods in Windows 95 and 98, you can use the AEGIS dot1x client or other comparable client software)

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.
- ◆ **Identity Profile User Name** – The dot1x supplicant user name. (Range: 1-8 characters)

The global supplicant user name and password are used to identify this switch as a supplicant when responding to an MD5 challenge from the authenticator. These parameters must be set when this switch passes client authentication requests to another authenticator on the network (see [“Configuring Port Supplicant Settings for 802.1X” on page 356](#)).
- ◆ **Set Password** – Allows the dot1x supplicant password to be entered.
- ◆ **Identity Profile Password** – The dot1x supplicant password used to identify this switch as a supplicant when responding to an MD5 challenge from the authenticator. (Range: 1-8 characters)
- ◆ **Confirm Profile Password** – This field is used to confirm the dot1x supplicant password.

- ◆ **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. Then set the user name and password to use when the switch responds an MD5 challenge from the authentication server.
4. Click Apply

Figure 221: Configuring Global Settings for 802.1X Port Authentication

Security > Port Authentication

Step: 1. Configure Global

System Authentication Control ☐ Enabled

EAPOL Pass Through ☐ Enabled

Identity Profile User Name

☒ Set Password

Identity Profile Password

Confirm Profile Password

Apply Revert

Default [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 – Authenticator) 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.
- ◆ When devices attached to a port must submit requests to another authenticator on the network, configure the Identity Profile parameters on the Configure Global page (see [“Configuring 802.1X Global Settings” on page 350](#)) which identify this switch as a supplicant, and configure the supplicant

parameters for those ports which must authenticate clients through the remote authenticator (see [“Configuring Port Supplicant Settings for 802.1X” on page 356](#)).

- ◆ This switch can be configured to serve as the authenticator on selected ports by setting the Control Mode to Auto on this configuration page, and as a supplicant on other ports by the setting the control mode to Force-Authorized on this page and enabling the PAE supplicant on the Supplicant 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 276.](#))
- ◆ **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 162](#)) and mapped on each port (See ["Configuring Network Access for Ports" on page 298](#)).

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. Click Authenticator.
4. Modify the authentication settings for each port as required.
5. Click Apply

Figure 222: Configuring Interface Settings for 802.1X Port Authenticator

Security > Port Authentication

Step: 2. Configure Interface

Type: ☒ Authenticator ☐ Supplicant

Port: 1

Status: Disabled

Authorized: Yes

Control Mode: Force-Authorized

Operation Mode: Single-Host

Max 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: 0 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

Configuring Port Supplicant Settings for 802.1X

Use the Security > Port Authentication (Configure Interface – Supplicant) page to configure 802.1X port settings for supplicant requests issued from a port to an authenticator on another device. When 802.1X is enabled and the control mode is set to Force-Authorized (see [“Configuring Port Authenticator Settings for 802.1X” on page 351](#)), you need to configure the parameters for the client supplicant process if the client must be authenticated through another device in the network.

Command Usage

- ◆ When devices attached to a port must submit requests to another authenticator on the network, configure the Identity Profile parameters on the Configure Global page (see [“Configuring 802.1X Global Settings” on page 350](#)) which identify this switch as a supplicant, and configure the supplicant parameters for those ports which must authenticate clients through the remote authenticator on this configuration page. When PAE supplicant mode is enabled on a port, it will not respond to dot1x messages meant for an authenticator.
- ◆ This switch can be configured to serve as the authenticator on selected ports by setting the Control Mode to Auto on the Authenticator configuration page, and as a supplicant on other ports by the setting the control mode to Force-Authorized on that configuration page and enabling the PAE supplicant on the Supplicant configuration page.

Parameters

These parameters are displayed:

- ◆ **Port** – Port number.
- ◆ **PAE Supplicant** – Enables PAE supplicant mode. (Default: Disabled)
If the attached client must be authenticated through another device in the network, supplicant status must be enabled.
Supplicant status can only be enabled if PAE Control Mode is set to “Force-Authorized” on this port (see [“Configuring Port Authenticator Settings for 802.1X” on page 351](#)).
PAE supplicant status cannot be enabled if a port is a member of trunk or LACP is enabled on the port.
- ◆ **Authentication Period** – The time that a supplicant port waits for a response from the authenticator. (Range: 1-65535 seconds; Default: 30 seconds)
- ◆ **Held Period** – The time that a supplicant port waits before resending its credentials to find a new an authenticator. (Range: 1-65535 seconds; Default: 30 seconds)
- ◆ **Start Period** – The time that a supplicant port waits before resending an EAPOL start frame to the authenticator. (Range: 1-65535 seconds; Default: 30 seconds)

- ◆ **Maximum Start** – The maximum number of times that a port supplicant will send an EAP start frame to the client before assuming that the client is 802.1X unaware. (Range: 1-65535; Default: 3)
- ◆ **Authenticated** – Shows whether or not the supplicant has been authenticated.

Web Interface

To configure port authenticator settings for 802.1X:

1. Click Security, Port Authentication.
2. Select Configure Interface from the Step list.
3. Click Supplicant.
4. Modify the supplicant settings for each port as required.
5. Click Apply

Figure 223: Configuring Interface Settings for 802.1X Port Supplicant

The screenshot displays the 'Security > Port Authentication' web interface. At the top, the 'Step' dropdown is set to '2. Configure Interface'. Below this, the 'Type' section has two radio buttons: 'Authenticator' (unselected) and 'Supplicant' (selected). The 'Port' is set to '2' via a dropdown menu. The 'PAE Supplicant' checkbox is checked and labeled 'Enabled'. Below this are five input fields for time periods: 'Authentication Period (1-65535)' with value '30', 'Held Period (1-65535)' with value '60', 'Start Period (1-65535)' with value '30', and 'Maximum Start (1-65535)' with value '3'. The 'Authenticated' status is shown as 'No'. At the bottom right, there are 'Apply' and 'Revert' buttons.

Field	Value
Type	Supplicant
Port	2
PAE Supplicant	Enabled
Authentication Period (1-65535)	30
Held Period (1-65535)	60
Start Period (1-65535)	30
Maximum Start (1-65535)	3
Authenticated	No

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 22: 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 22: 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. Click Authenticator.

Figure 224: Showing Statistics for 802.1X Port Authenticator

Security > Port Authentication

Step: 3. Show Statistics

Type
☒ Authenticator
☐ Supplicant

Port
1

Port Authentication Authenticator Statistics

Rx EAPOL Start	11154	Rx EAP Resp/Id	2533664
Rx EAPOL Logoff	2115542	Rx EAP Resp/Oth	11123
Rx EAPOL Invalid	533	Rx EAP LenError	1
Rx EAPOL Total	1000	Tx EAP Req/Id	5222
Rx Last EAPOLVer	255	Tx EAP Req/Oth	1222
Rx Last EAPOLSrc	00-02-44-51-C2-90	Tx EAPOL Total	1

Refresh

To display port supplicant statistics for 802.1X:

1. Click Security, Port Authentication.
2. Select Show Statistics from the Step list.
3. Click Supplicant.

Figure 225: Showing Statistics for 802.1X Port Supplicant

Security > Port Authentication

Step: 3. Show Statistics

Type: ☐ Authenticator ☒ Supplicant

Port: 1

Port Authentication Supplicant Statistics

Rx EAPOL Invalid	11154	Rx EAP LenError	2533664
Rx EAPOL Total	2115542	Tx EAPOL Total	11123
Rx Last EAPOLVer	533	Tx EAPOL Start	1
Rx Last EAPOLSrc	1000	Tx EAPOL Logoff	5222
Rx EAP Resp/Id	255	Tx EAP Req/Id	1222
Rx EAP Resp/Oth	00-02-44-51-C2-90	Tx EAP Req/Oth	1

Refresh

DoS Protection

Use the Security > DoS Protection page to protect against denial-of-service (DoS) attacks. A DoS attack is an attempt to block the services provided by a computer or network resource. This kind of attack tries to prevent an Internet site or service from functioning efficiently or at all. In general, DoS attacks are implemented by either forcing the target to reset, to consume most of its resources so that it can no longer provide its intended service, or to obstruct the communication media between the intended users and the target so that they can no longer communicate adequately. This section describes how to protect against DoS attacks.

Parameters

These parameters are displayed:

- ◆ **Echo/Chargen Attack** – Attacks in which the echo service repeats anything sent to it, and the chargen (character generator) service generates a continuous stream of data. When used together, they create an infinite loop and result in a denial-of-service. (Default: Disabled)
- ◆ **Echo/Chargen Attack Rate** – Maximum allowed rate. (Range: 64-2000 kbits/second; Default: 1000 kbits/second)

- ◆ **Smurf Attack** – Attacks in which a perpetrator generates a large amount of spoofed ICMP Echo Request traffic to the broadcast destination IP address (255.255.255.255), all of which uses a spoofed source address of the intended victim. The victim should crash due to the many interrupts required to send ICMP Echo response packets. (Default: Enabled)
- ◆ **TCP Flooding Attack** – Attacks in which a perpetrator sends a succession of TCP SYN requests (with or without a spoofed-Source IP) to a target and never returns ACK packets. These half-open connections will bind resources on the target, and no new connections can be made, resulting in a denial of service. (Default: Disabled)
- ◆ **TCP Flooding Attack Rate** – Maximum allowed rate. (Range: 64-2000 kbits/second; Default: 1000 kbits/second)
- ◆ **TCP Null Scan** – A TCP NULL scan message is used to identify listening TCP ports. The scan uses a series of strangely configured TCP packets which contain a sequence number of 0 and no flags. If the target's TCP port is closed, the target replies with a TCP RST (reset) packet. If the target TCP port is open, it simply discards the TCP NULL scan. (Default: Enabled)
- ◆ **TCP-SYN/FIN Scan** – A TCP SYN/FIN scan message is used to identify listening TCP ports. The scan uses a series of strangely configured TCP packets which contain SYN (synchronize) and FIN (finish) flags. If the target's TCP port is closed, the target replies with a TCP RST (reset) packet. If the target TCP port is open, it simply discards the TCP SYN FIN scan. (Default: Enabled)
- ◆ **TCP Xmas Scan** – A so-called TCP XMAS scan message is used to identify listening TCP ports. This scan uses a series of strangely configured TCP packets which contain a sequence number of 0 and the URG, PSH and FIN flags. If the target's TCP port is closed, the target replies with a TCP RST packet. If the target TCP port is open, it simply discards the TCP XMAS scan. (Default: Enabled)
- ◆ **TCP/UDP Packets with Port 0** – Protects against DoS attacks in which the TCP or UDP source port or destination port is set to zero. This technique may be used as a form of DoS attack, or it may just indicate a problem with the source device. When this command is enabled, the switch will drop these packets. (Default: Enabled)



Note: Due to a chip limitation, this command does not work when the TCP source port/destination port or UDP source port/destination port are both set to zero.

- ◆ **UDP Flooding Attack** – Attacks in which a perpetrator sends a large number of UDP packets (with or without a spoofed-Source IP) to random ports on a remote host. The target will determine that application is listening at that port, and reply with an ICMP Destination Unreachable packet. It will be forced to send many ICMP packets, eventually leading it to be unreachable by other clients. (Default: Disabled)

- ◆ **UDP Flooding Attack Rate** – Maximum allowed rate. (Range: 64-2000 kbits/second; Default: 1000 kbits/second)
- ◆ **WinNuke Attack** – Attacks in which affected the Microsoft Windows 3.1x/95/NT operating systems. In this type of attack, the perpetrator sends the string of OOB out-of-band (OOB) packets contained a TCP URG flag to the target computer on TCP port 139 (NetBIOS), casing it to lock up and display a “Blue Screen of Death.” This did not cause any damage to, or change data on, the computer’s hard disk, but any unsaved data would be lost. Microsoft made patches to prevent the WinNuke attack, but the OOB packets. (Default: Disabled)
- ◆ **WinNuke Attack Rate** – Maximum allowed rate. (Range: 64-2000 kbits/second; Default: 1000 kbits/second)

Web Interface

To protect against DoS attacks:

1. Click Security, DoS Protection.
2. Enable protection for specific DoS attacks, and set the maximum allowed rate as required.
3. Click Apply

Figure 226: Protecting Against DoS Attacks

The screenshot displays the 'Security > DoS Protection' configuration page. It features a list of attack types on the left and their corresponding status and rate settings on the right. The settings are as follows:

Attack Type	Status	Rate (kbps)
Echo/Chargen Attack	☐ Enabled	
Echo/Chargen Attack Rate (64-2000)		1000
Smurf Attack	☒ Enabled	
TCP Flooding Attack	☐ Enabled	
TCP Flooding Attack Rate (64-2000)		1000
TCP Null Scan	☒ Enabled	
TCP SYN/FIN Scan	☒ Enabled	
TCP Xmas Scan	☒ Enabled	
TCP/UDP Packets With Port 0	☒ Enabled	
UDP Flooding Attack	☐ Enabled	
UDP Flooding Attack Rate (64-2000)		1000
WinNuke Attack	☐ Enabled	
WinNuke Attack Rate (64-2000)		1000

At the bottom right of the interface are two buttons: 'Apply' and 'Revert'.

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 375](#)). IP 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 IPv4 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

Filter Type

- ◆ 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 375](#)), 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 377](#)), IP 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 IP source guard binding, the packet will be forwarded.

- If DHCP snooping is enabled, IP 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 IP source guard binding, or dynamic DHCP snooping binding, the packet will be forwarded.
- If IP source guard is enabled on an interface for which IP source bindings have not yet been configured (neither by static configuration in the IP 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 IP 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 375](#)) and static entries set by IP source guard (see [“Configuring Static Bindings for IPv4 Source Guard” on page 365](#)).

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 227: Setting the Filter Type for IPv4 Source Guard

Port	Filter Type	Filter Table	ACL Table Max Binding Entry (1-16)	MAC Table Max Binding Entry (1-1024)
1	SIP	ACL	5	1024
2	SIP-MAC	ACL	5	1024
3	DISABLED	ACL	5	1024
4	DISABLED	ACL	5	1024
5	DISABLED	ACL	5	1024

Configuring Static Bindings for IPv4 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:

Add – Configure ACL Table

- ◆ **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.

Add – Configure MAC Table

- ◆ **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: 1-28/52)

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 228: 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 Configure ACL Table or Configure MAC Table from the Step list.
3. Select Show from the Action list.

Figure 229: Displaying Static Bindings for IPv4 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
☐	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.

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 230: Showing the IPv4 Source Guard Binding Table

Security > IP Source Guard > Dynamic Binding

Query by:

☐ Port

☐ VLAN

☐ 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 IP 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: 1-28/52)
- ◆ **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 371).

- 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 231: Setting the Filter Type for IPv6 Source Guard

Security > IPv6 Source Guard > Port Configuration		
Port Configuration List Total: 28		
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 Binding 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 232: 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 233: Displaying Static Bindings for IPv6 Source Guard

Security > IPv6 Source Guard > Static Binding

Action: Show

Static Binding List Total: 3

	VLAN	MAC Address	Interface	IPv6 Address	Type
☐	1	00-10-B5-F4-00-01	Eth 1/2	2001:DB8:2222:7272::26	STA
☐	1	00-10-B5-F4-00-02	Eth 1/4	2001:DB8:2222:7272::56	DHCP
☐	2	00-10-B5-F4-00-03	Eth 1/7	2001:DB8:2222:7272::36	ND

Delete Revert

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 234: Showing the IPv6 Source Guard Binding Table



Security > IPv6 Source Guard > Dynamic Binding

Query by:

☐ Port

☐ VLAN

☐ MAC Address

☐ IPv6 Address

Query

Dynamic Binding List Total: 3

VLAN	MAC Address	Interface	IPv6 Address	Type
1	00-10-B5-F4-00-01	Eth 1/2	2001::de	ND
1	00-10-B5-F4-00-02	Eth 1/4	2001::ee	ND
2	00-10-B5-F4-00-03	Pch 1	2001::fe	ND

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 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: Enabled)
- ◆ **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 snooping information option.
4. Click Apply

Figure 235: Configuring Global Settings for DHCP Snooping

The screenshot shows the configuration interface for DHCP Snooping. The breadcrumb path is IP Service > DHCP > Snooping. The 'Step' dropdown is set to '1. Configure Global'. The 'General' section includes: 'DHCP Snooping Status' (checkbox, Enabled), 'DHCP Snooping MAC-Address Verification' (checkbox, checked, Enabled), and 'DHCP Snooping Rate Limit' (checkbox, 0 pkt/sec). The 'Information' section includes: 'DHCP Snooping Information Option Status' (checkbox, Enabled), 'DHCP Snooping Information Option Sub-option Format' (dropdown, Extra Subtype Included), 'DHCP Snooping Information Option Remote ID' (dropdown, MAC Address (Hex Encoded)), 'DHCP Snooping Information Option Remote ID TR101 VLAN Field' (checkbox, checked, Enabled), 'DHCP Snooping Information Option TR101 Board ID' (checkbox, none), and 'DHCP Snooping Information Option Policy' (dropdown, Replace). 'Apply' and 'Revert' buttons are at the bottom right.

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 236: Configuring DHCP Snooping on a VLAN

The screenshot shows a web interface for configuring DHCP Snooping. The breadcrumb trail is "IP Service > DHCP > Snooping". Below this, there is a "Step:" dropdown menu currently set to "2. Configure VLAN". Underneath, there is a "VLAN" dropdown menu set to "1". Below that, the "DHCP Snooping Status" is shown as "Enabled" with a green checkmark icon. At the bottom right, there are two buttons: "Apply" and "Revert".

Configuring Ports for DHCP Snooping

Use the IP Service > DHCP > Snooping (Configure Interface) page to configure switch ports 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 ports within the VLAN.
- ◆ When an untrusted port is changed to a trusted port, all the dynamic DHCP snooping bindings associated with this port are removed.
- ◆ Set all ports connected to DHCP servers within the local network or fire wall to trusted state. Set all other ports outside the local network or fire wall to untrusted state.

Parameters

These parameters are displayed:

- ◆ **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 237: Configuring the Port Mode for DHCP Snooping

IP Service > DHCP > Snooping

Step: 3. Configure Interface

Interface ☒ Port ☐ Trunk

DHCP Snooping Port List Total: 28

Port	Trust Status	Circuit ID	
		Mode	Value
1	☐ Enabled	VLAN-Unit-Port	
2	☐ Enabled	VLAN-Unit-Port	
3	☐ Enabled	VLAN-Unit-Port	
4	☐ Enabled	VLAN-Unit-Port	
5	☒ Enabled	String	RD-934

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-DHCPSPNP** – 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 238: Displaying the Binding Table for DHCP Snooping



The screenshot shows the web interface for IP Service > DHCP > Snooping. The 'Step' dropdown is set to '4. Show Information'. Below the dropdown is the 'DHCP Snooping Binding List' table, which has a 'Total: 6' label. The table contains 6 rows of data. Below the table are two buttons: 'Store' and 'Clear', each with a description of its function.

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-DHCPSPNP	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-DHCPSPNP	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-DHCPSPNP	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.

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.
- ◆ **Switch Clustering** – Configures centralized management by a single unit over a group of switches connected to the same local network.
- ◆ **Ethernet Ring Protection Switching (ERPS)** – Configures a protection switching mechanism and protocol for Ethernet layer network rings.
- ◆ **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.
- ◆ **Operation, Administration and Maintenance (OAM)** – Provides remote management tools required to monitor and maintain the links to subscriber CPEs (Customer Premise Equipment).
- ◆ **UniDirectional Link Detection (UDLD)** – Detects general loopback conditions caused by hardware problems or faulty protocol settings.

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 23: 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 239: Configuring Settings for System Memory Logs

Administration > Log > System

Step: 1. Configure Global

System Log Status ☒ Enabled

Flash Level 3 - Error

RAM Level 7 - Debugging

Note: The Flash Level must be equal to or less than the RAM Level.

Apply 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 240: 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)

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 241: Configuring Settings for Remote Logging of Error Messages

Sending Simple Mail Transfer Protocol Alerts

Use the Administration > Log > SMTP page to alert system administrators of problems by sending SMTP (Simple Mail Transfer Protocol) email messages when triggered by logging events of a specified level. The messages are sent to specified SMTP servers on the network and can be retrieved using POP or IMAP clients.

Parameters

These parameters are displayed:

- ◆ **SMTP Status** – Enables/disables the SMTP function. (Default: Enabled)
- ◆ **Severity** – Sets the syslog severity threshold level (see table on [page 384](#)) used to trigger alert messages. All events at this level or higher will be sent to the configured email recipients. For example, using Level 7 will report all events from level 7 to level 0. (Default: Level 7)
- ◆ **Email Source Address** – Sets the email address used for the “From” field in alert messages. You may use a symbolic email address that identifies the switch, or the address of an administrator responsible for the switch. (Range: 1-41 characters)

- ◆ **Email Destination Address** – Specifies the email recipients of alert messages. You can specify up to five recipients.
- ◆ **Server IP Address** – Specifies a list of up to three recipient SMTP servers. IPv4 or IPv6 addresses may be specified. The switch attempts to connect to the listed servers in sequential order if the first server fails to respond.

For host name-to-IP address translation to function properly, host name lookup must be enabled ([“Configuring General DNS Service Parameters” on page 637](#)), and one or more DNS servers specified (see [“Configuring a List of Name Servers” on page 640](#), or [“Configuring Static DNS Host to Address Entries” on page 641](#)).

Web Interface

To configure SMTP alert messages:

1. Click Administration, Log, SMTP.
2. Enable SMTP, specify a source email address, and select the minimum severity level. Specify the source and destination email addresses, and one or more SMTP servers.
3. Click Apply.

Figure 242: Configuring SMTP Alert Messages

The screenshot shows the 'Administration > Log > SMTP' configuration page. It contains the following fields and controls:

- SMTP Status:** A checkbox labeled 'Enabled' which is checked.
- Severity:** A dropdown menu currently showing '3 - Error'.
- E-mail Source Address:** A text input field containing 'big-wheels@matel.com'.
- E-mail Destination Address 1:** A text input field containing 'chris@matel.com'.
- E-mail Destination Address 2:** An empty text input field.
- E-mail Destination Address 3:** An empty text input field.
- E-mail Destination Address 4:** An empty text input field.
- E-mail Destination Address 5:** An empty text input field.
- Server IP Address 1:** A text input field containing '192.168.1.4'.
- Server IP Address 2:** An empty text input field.
- Server IP Address 3:** An empty text input field.

At the bottom right of the form are two buttons: 'Apply' and 'Revert'.

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.

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 243: 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 has a checkbox labeled 'Enabled' which is checked. Underneath, several timing attributes are listed with input fields and units:

- 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

A note at the bottom states: 'Note: The Transmission Interval must be greater than or equal to 4 times the Delay Interval.' At the bottom right, there are 'Apply' and 'Revert' buttons.

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 432](#).

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. (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 80](#). (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 (see [“Protocol VLANs” on page 181](#)). (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 159](#)). (Default: Enabled)
- **VLAN Name** – The name of all VLANs to which this interface has been assigned (see [“IEEE 802.1Q VLANs” on page 159](#)). (Default: Enabled)
- **Port and Protocol VLAN ID** – The port-based protocol VLANs configured on this interface (see [“Protocol VLANs” on page 181](#)). (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 82](#) 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 244: Configuring LLDP Interface Attributes

The screenshot shows the 'Administration > LLDP' configuration page. At the top, 'Step: 2. Configure Interface' and 'Action: Configure General' are selected. The 'Interface' section has 'Port 1' selected. 'Admin Status' is set to 'Tx Rx'. 'SNMP Notification' and 'MED Notification' are both unchecked. Under 'Basic Optional TLVs', 'Management Address', 'Port Description', 'System Capabilities', 'System Description', and 'System Name' are unchecked. 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', 'MAC/PHY Configuration/Status', 'PoE', and 'MED TLVs' (Capabilities, Extended Power, Inventory, Location, Network Policy) are unchecked. The 'MED-Location Civic Address' section shows 'Country' as 'US' and 'Device entry refers to' as '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'. 'Apply' and 'Revert' buttons are at the bottom right.

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 24: 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
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 245: Configuring the Civic Address for an LLDP Interface

Administration > LLDP

Step: 2. Configure Interface Action: Add CA-Type

Interface ☒ Port 1 ☐ Trunk

CA-Type (0-255) 1

CA-Value California

Apply Revert

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 246: Showing the Civic Address for an LLDP Interface

Administration > LLDP

Step: 2. Configure Interface Action: Show CA-Type

Interface ☒ Port 1 ☐ Trunk

LLDP MED-Location Civic Address Type List Total: 1

	CA-Type	CA-Value
☐	1	California

Delete Revert

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 25: 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 80](#)).
- ◆ **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 26: 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

Table 26: System Capabilities (Continued)

ID Basis	Reference
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 27: 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)

Table 27: Port ID Subtype (Continued)

ID Basis	Reference
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 247: Displaying Local Device Information for LLDP (General)

Administration > LLDP


Step: 3. Show Local Device Information ▼

☒ General
 ☐ Port
 ☐ Port Details
 ☐ Trunk
 ☐ Trunk Details

LLDP Local Device Information

Chassis Type	MAC Address
Chassis ID	00-E0-0C-00-00-FD
System Name	
System Description	GTL-2891
System Capabilities Supported	Bridge, Router
System Capabilities Enabled	Bridge, Router
Management Address	192.168.2.32 (IPv4)

Figure 248: Displaying Local Device Information for LLDP (Port)



Administration > LLDP


Step: 3. Show Local Device Information

General Port Port Details Trunk Trunk Details

LLDP Local Device Port List Total: 28

Port	Port Description	Port ID
1	Ethernet Port on unit 1, port 1	00-00-E8-94-40-01
2	Ethernet Port on unit 1, port 2	00-00-E8-94-40-02
3	Ethernet Port on unit 1, port 3	00-00-E8-94-40-03
4	Ethernet Port on unit 1, port 4	00-00-E8-94-40-04
5	Ethernet Port on unit 1, port 5	00-00-E8-94-40-05

Figure 249: Displaying Local Device Information for LLDP (Port Details)



Administration > LLDP

Step: 3. Show Local Device Information

General Port Port Details Trunk Trunk Details

Port 1

LLDP Local Port Information Details

Local Port	1
Port ID Type	MAC Address
Port ID	00-00-E8-94-40-01
Port Description	Ethernet Port on unit 1, port 1
MED Capability	LLDP-MED Capabilities, Network Policy, Location Identification, Inventory

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 25, "Chassis ID Subtype," on page 397.](#))
- ◆ **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 27, "Port ID Subtype," on page 398.](#)
- ◆ **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 26, "System Capabilities," on page 397.](#))
- ◆ **System Capabilities Enabled** – The primary function(s) of the system which are currently enabled. (See [Table 26, "System Capabilities," on page 397.](#))
- ◆ **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 28: 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.
- ◆ **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 Policy⁸

- ◆ **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.

8. These fields are only displayed for end-node devices advertising LLDP-MED TLVs.

- ◆ **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 are 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:

9. Location Configuration Information

- **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.
- ◆ **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 250: Displaying Remote Device Information for LLDP (Port)

Administration > LLDP

Step: 4. Show Remote Device Information

☒ Port ☐ Port Details ☐ Trunk ☐ Trunk Details

LLDP Remote Device Port List Total: 1

Local Port	Chassis ID	Port ID	System Name
1	70-72-CF-32-DD-FD	70-72-CF-32-DD-FF	RD93

Administration > LLDP

Step: 4. Show Remote Device Information

☐ Port ☒ Port Details ☐ Trunk ☐ Trunk Details

Port 1

Remote Index 2

Query

Figure 251: Displaying Remote Device Information for LLDP (Port Details)

Administration > LLDP

Step: 4. Show Remote Device Information

Port
Port Details
Trunk
Trunk Details

Port 23
Remote Index 4

Query

LLDP Remote Device Port Information

Local Port	23	Port Type	MAC Address
Chassis Type	MAC Address	Port Description	Ethernet Port on unit 1, port 1
Chassis ID	00-E0-00-00-00-01	Port ID	00-E0-00-00-00-02
System Name		System Capabilities Supported	Bridge
System Description	ES3528MV2	System Capabilities Enabled	Bridge

Management Address List Total: 1

Address	Address Type
192.168.0.3	IPv4 Address

802.1 Extension Information

Remote Port VID 1

Remote Port-Protocol VLAN List Total: 1

VLAN	Support	Status
1	Yes	Enabled

Remote VLAN Name List Total: 1

VLAN	Name
1	DefaultVlan

Remote Protocol Identity List Total: 1

Remote Protocol Identity (Hex)
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	0000	Remote Port MAU Type	6

802.3 Extension Power Information

Remote Power Class	PSE	Remote Power MDI Supported	Yes
Remote Power MDI Status	Enabled	Remote Power Pair Controlable	No
Remote Power Pairs	Spare	Remote Power Classification	Class1

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	1518
-----------------------	------

Additional information displayed by an end-point device which advertises LLDP-MED TLVs is shown in the following figure.

Figure 252: 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

Location Identification

Location Data Format	Civic Address LCI		
Country Code	US	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	1.2.2.1
Software Revision	1.2.2.1	Serial Number	LN10230092
Manufacture Name		Model Name	L
Asset ID			

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 253: Displaying LLDP Device Statistics (General)

LLDP Device Statistics	
Neighbor Entries List Last Updated	2 sec
New Neighbor Entries Count	20
Neighbor Entries Deleted Count	20
Neighbor Entries Dropped Count	0
Neighbor Entries Age-out Count	20

Figure 254: Displaying LLDP Device Statistics (Port)

LLDP Device Port Statistics			
Frames Discarded	0	TLVs Unrecognized	0
Frames Invalid	0	TLVs Discarded	0
Frames Received	97	Neighbor Ageouts	0
Frames Sent	104		

Refresh

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 29: 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)

- ◆ **Link-up and Link-down Traps**¹⁰ – Issues a notification message whenever a port link is established or broken. (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

Figure 255: Configuring Global Settings for SNMP

The screenshot shows a web interface for configuring SNMP. The breadcrumb is "Administration > SNMP". Below it is a "Step:" dropdown menu with "1. Configure Global" selected. The main configuration area has three rows, each with a label and a checkbox: "Agent Status" with a checked checkbox and the word "Enabled"; "Authentication Traps" with a checked checkbox and the word "Enabled"; and "Link-up and Link-down Traps" with a checked checkbox and the word "Enabled". At the bottom right of the configuration area are two buttons: "Apply" and "Revert".

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".

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 417).

- ◆ **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 256: Configuring the Local Engine ID for SNMP

The screenshot shows a web interface for configuring SNMP. At the top, it says "Administration > SNMP". Below this, there are two dropdown menus: "Step:" with "2. Configure Engine" selected, and "Action:" with "Set Engine ID" selected. Below these, there is a text input field for "Engine ID" containing the hexadecimal string "800001030300000c0000fd0000". Below the input field, there is a label "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 429.](#))

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 or IPv6 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 257: Configuring a Remote Engine ID for SNMP

The screenshot shows the 'Administration > SNMP' web interface. At the top, there is a breadcrumb trail. Below it, there are two dropdown menus: 'Step: 2. Configure Engine' and 'Action: Add Remote Engine'. The main form has two input fields: 'Remote Engine ID' with the value '5432100000' and 'Remote IP Host' with the value '192.168.1.19'. At the bottom right, there are two buttons: 'Apply' and '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 258: Showing Remote Engine IDs for SNMP

The screenshot shows the 'Administration > SNMP' web interface. At the top, there is a breadcrumb trail. Below it, there are two dropdown menus: 'Step: 2. Configure Engine' and 'Action: Show Remote Engine'. The main form displays a table titled 'SNMPv3 Remote Engine List' with a 'Total: 1' indicator. The table has two columns: 'Remote Engine ID' and 'Remote IP Host'. The first row shows the values '5432100000' and '192.168.1.19'. At the bottom right, there are two buttons: 'Delete' and 'Revert'.

Remote Engine ID	Remote IP Host
5432100000	192.168.1.19

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** – The name of the SNMP view. (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. (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 259: 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 260: Showing SNMP Views

Administration > SNMP

Step: 3. Configure View Action: Show View

SNMPv3 View List Total: 2

☐	View Name
☐	ifEntry.a
☐	defaultview

Delete Revert

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 261: Adding an OID Subtree to an SNMP View

Administration > SNMP

Step: 3. Configure View Action: Add OID Subtree

View Name: ifEntry.a

OID Subtree: 1.3.6.1.2.1.2.2.1.2.*

Type: Included

Apply Revert

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 262: Showing the OID Subtree Configured for SNMP Views

Administration > SNMP

Step: 3. Configure View Action: Show OID Subtree

View Name: ifEntry.a

SNMPv3 View OID Subtree List Total: 2

	OID Subtree	Type
☐	1.3.6.1.2.1.2.2.1.1.*	Included
☐	1.3.6.1.2.1.2.2.1.2.*	Included

Delete Revert

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. (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 30: 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.

Table 30: Supported Notification Messages (Continued)

Model	Level	Group
<i>Private Traps</i>		
swPowerStatusChangeTrap	1.3.6.1.4.1.22426.10.1.41.2.1.0.1	This trap is sent when the power state changes.
swPortSecurityTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.36	This trap is sent when the port is being intruded. This trap will only be sent when the portSecActionTrap is enabled.
swIpFilterRejectTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.40	This trap is sent when an incorrect IP address is rejected by the IP Filter.
swAtcBcastStormAlarmFireTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.70	When broadcast traffic is detected as a storm, this trap is fired.
swAtcBcastStormAlarmClearTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.71	When a broadcast storm is detected as normal traffic, this trap is fired.
swAtcBcastStormTcApplyTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.72	When ATC is activated, this trap is fired.
swAtcBcastStormTcReleaseTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.73	When ATC is released, this trap is fired.
swAtcMcastStormAlarmFireTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.74	When multicast traffic is detected as the storm, this trap is fired.
swAtcMcastStormAlarmClearTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.75	When multicast storm is detected as normal traffic, this trap is fired.
swAtcMcastStormTcApplyTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.76	When ATC is activated, this trap is fired.
swAtcMcastStormTcReleaseTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.77	When ATC is released, this trap is fired.
stpBpduGuardPortShutdownTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.91	This trap will be sent when an interface is shut down because of BPDU guard.
swLoopbackDetectionTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.95	This trap is sent when loopback BPDUs have been detected.
networkAccessPortLinkDetectionTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.96	This trap is sent when a networkAccessPortLinkDetection event is triggered.
dot1agCfmMepUpTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.97	This trap is sent when a new remote MEP is discovered.
dot1agCfmMepDownTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.98	This trap is sent when port status or interface status TLV received from remote MEP indicates it is not up.
dot1agCfmConfigFailTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.99	This trap is sent when a MEP receives a CCM with MPID which already exists on the same MA in this switch.
dot1agCfmLoopFindTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.100	This trap is sent when a MEP receives its own CCMs.
dot1agCfmMepUnknownTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.101	This trap is sent when a CCM is received from an unexpected MEP.
dot1agCfmMepMissingTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.102	This trap is sent when the cross-check enable timer expires and no CCMs were received from an expected (configured) MEP.
dot1agCfmMaUpTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.103	This trap is sent when all expected remote MEPs are up.
autoUpgradeTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.104	This trap is sent when auto upgrade is executed.
swCpuUtiRisingNotification	1.3.6.1.4.1.22426.10.11.41.2.1.0.107	This notification indicates that the CPU utilization has risen from cpuUtiFallingThreshold to cpuUtiRisingThreshold.

Table 30: Supported Notification Messages (Continued)

Model	Level	Group
swCpuUtiFallingNotification	1.3.6.1.4.1.22426.10.11.41.2.1.0.108	This notification indicates that the CPU utilization has fallen from <code>cpuUtiRisingThreshold</code> to <code>cpuUtiFallingThreshold</code> .
swMemoryUtiRisingThreshold Notification	1.3.6.1.4.1.22426.10.11.41.2.1.0.109	This notification indicates that the memory utilization has risen from <code>memoryUtiFallingThreshold</code> to <code>memoryUtiRisingThreshold</code> .
swMemoryUtiFallingThreshold Notification	1.3.6.1.4.1.22426.10.11.41.2.1.0.110	This notification indicates that the memory utilization has fallen from <code>memoryUtiRisingThreshold</code> to <code>memoryUtiFallingThreshold</code> .
dhcpRogueServerAttackTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.114	This trap is sent when receiving a DHCP packet from a rogue server.
macNotificationTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.138	This trap is sent when there are changes of the dynamic MAC addresses on the switch.
lbdDetectionTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.124	This trap is sent when a loopback condition is detected by LBD.
lbdRecoveryTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.142	This trap is sent when a recovery is done by LBD.
sfpThresholdAlarmWarnTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.189	This trap is sent when the SFP's monitored value is not within alarm/warning thresholds.
udldPortShutdownTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.192	This trap is sent when the port is shut down by UDLD.
userAuthenticationFailureTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.199	This trap will be triggered if authentication fails.
userAuthenticationSuccessTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.200	This trap will be triggered if authentication is successful.
loginTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.201	This trap is sent when user logs in.
logoutTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.202	This trap is sent when user logs out.
fileCopyTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.208	This trap is sent when file copy is executed. If the copy action is triggered by the system, the login user information (<code>trapVarLoginUserName/</code> <code>trapVarSessionType/</code> <code>trapVarLoginInetAddressTypes/</code> <code>trapVarLoginInetAddress</code>) will be a null value.
userauthCreateUserTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.209	This trap is sent when a user account is created.
userauthDeleteUserTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.210	This trap is sent when a user account is deleted.
userauthModifyUserPrivilegeTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.211	This trap is sent when user privilege is modified.
cpuGuardControlTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.213	This trap is sent when CPU utilization rises above the high-watermark the first time or when CPU utilization rises from below the low-watermark to above the high-watermark.
cpuGuardReleaseTrap	1.3.6.1.4.1.22426.10.11.41.2.1.0.214	This trap is sent when CPU utilization falls from above the high-watermark to below the low-watermark.

* 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 263: Creating an SNMP Group

Administration > SNMP

Step: 4. Configure Group Action: Add

Group Name: secure-users

Security Model: v3

Security Level: authPriv

Read View: ☒ ifEntry.a ☐

Write View: ☒ ifEntry.a ☐

Notify View: ☒ ifEntry.a ☐

Apply Revert

To show SNMP groups:

1. Click Administration, SNMP.
2. Select Configure Group from the Step list.
3. Select Show from the Action list.

Figure 264: Showing SNMP Groups

Administration > SNMP

Step: 4. Configure Group Action: Show

SNMPv3 Group List Total: 4

	Group Name	Model	Level	Read View	Write View	Notify View
☐	public	v1	noAuthNoPriv	defaultview	No writeview specified	No notifyview specified
☐	public	v2c	noAuthNoPriv	defaultview	No writeview specified	No notifyview specified
☐	private	v1	noAuthNoPriv	defaultview	defaultview	No notifyview specified
☐	private	v2c	noAuthNoPriv	defaultview	defaultview	No notifyview specified

Delete Revert

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 265: Setting Community Access Strings

Administration > SNMP

Step: 5. Configure User Action: Add Community

Community String spiderman

Access Mode Read/Write

Apply Revert

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 266: Showing Community Access Strings

Administration > SNMP		
Step:	5. Configure User	Action: Show Community
SNMP Community String List Total: 3		
	Community String	Access Mode
☐	public	Read-Only
☐	private	Read/Write
☐	spiderman	Read/Write

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.
(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. (Range: 8-32 characters)
- ◆ **Privacy Protocol** – The encryption algorithm use for data privacy; only 56-bit DES is currently available.
- ◆ **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 267: Configuring Local SNMPv3 Users

The screenshot shows the 'Administration > SNMP' configuration page. At the top, the 'Step' is set to '5. Configure User' and the 'Action' is 'Add SNMPv3 Local User'. The form is divided into several sections:

- SNMPv3 User**
 - User Name:
 - Group Name: ☐ public ☒ r&d
 - Security Model:
 - Security Level:
- User Authentication**
 - Authentication Protocol:
 - Authentication Password:
- Data Privacy**
 - Privacy Protocol:
 - Privacy Password:

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

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 268: Showing Local SNMPv3 Users

The screenshot shows the 'Administration > SNMP' configuration page with the 'Action' set to 'Show SNMPv3 Local User'. Below the configuration fields, there is a table titled 'SNMPv3 Local User List' with a 'Total: 1' indicator.

	User Name	Group Name	Model	Level	Authentication	Privacy
☐	chris	r&d	v3	authPriv	MD5	DES56

At the bottom right, there are 'Delete' and 'Revert' buttons.

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 269: Changing a Local SNMPv3 User Group

The screenshot shows a web interface for configuring SNMPv3 users. The breadcrumb is 'Administration > SNMP'. Below it, there are two dropdown menus: 'Step: 5. Configure User' and 'Action: Change SNMPv3 Local User Group'. The main configuration area has two fields: 'User Name' with a dropdown menu showing 'chris', and 'Group Name' with two radio buttons. The 'bart' radio button is selected, and the 'public' radio button is unselected. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

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 432](#) and [“Specifying a Remote Engine ID” on page 415.](#))

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 or IPv6 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 use for data privacy; only 56-bit DES is currently available.
- ◆ **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 270: Configuring Remote SNMPv3 Users

Administration > SNMP

Step: 5. Configure User Action: Add SNMPv3 Remote User

SNMPv3 User

User Name: mark

Group Name: ☐ public ☒ r&d

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

Apply Revert

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 271: Showing Remote SNMPv3 Users

Administration > SNMP

Step: 5. Configure User Action: Show SNMPv3 Remote User

SNMPv3 Remote User List Total: 1

	User Name	Group Name	Engine ID	Model	Level	Authentication	Privacy
☐	mark	r&d	5432100000	v3	authPriv	MD5	DES56

Delete Revert

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 413](#)).
2. Create a view with the required notification messages ([page 417](#)).
3. Configure the group (matching the community string specified on the Configure Trap - Add page) to include the required notify view ([page 419](#)).
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 413](#)).
2. Create a remote SNMPv3 user to use in the message exchange process ([page 426](#)). 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 417](#)).
4. Create a group that includes the required notify view ([page 419](#)).
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 as 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 as 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 426](#)), 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 429](#)), 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 272: Configuring Trap Managers (SNMPv1)

Administration > SNMP

Step: 6. Configure Trap Action: Add

IP Address 192.168.0.3

Version v1

Community String private

UDP Port (1-65535) 162

Apply Revert

Figure 273: Configuring Trap Managers (SNMPv2c)

Administration > SNMP

Step: 6. Configure Trap Action: Add

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)

Apply Revert

Figure 274: Configuring Trap Managers (SNMPv3)

Administration > SNMP

Step: 6. Configure Trap Action: Add

IP Address 192.168.2.9

Version v3

Notification Type Inform

Timeout (0-2147483647) centiseconds

Retry Times (0-255)

Remote User Name

UDP Port (1-65535)


Security Level authPriv

Apply Revert

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 275: Showing Trap Managers



	IP Address	Version	Community String/User Name	UDP Port	Security Level	Timeout	Retry Times
☐	192.168.0.4	v3	steve	162	noAuthNoPriv		
☐	192.168.0.5	v3	bobby	162	noAuthNoPriv		
☐	192.168.0.6	v3	betty	162	authNoPriv		
☐	192.168.2.9	v2c	venus	162		1600	5
☐	192.168.5.8	v3	margaret	162	authPriv	1600	5

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 87](#). 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 432](#), a default notify filter will be created.

Parameters

These parameters are displayed:

- ◆ **IP Address** – The IPv4 or IPv6 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 276: Creating SNMP Notification Logs

The screenshot shows a web interface titled "Administration > SNMP". Below the title, there are two dropdown menus: "Step:" with "7. Configure Notify Filter" selected, and "Action:" with "Add" selected. Below these, there are two text input fields: "IP Address" with the value "192.168.0.99" and "Filter Profile Name" with the value "R&D". At the bottom right, there are two buttons: "Apply" and "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 277: Showing SNMP Notification Logs



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.

Web Interface

To show SNMP statistics:

1. Click Administration, SNMP.
2. Select Show Statistics from the Step list.

Figure 278: 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		
Refresh			

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 falling 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 279: Configuring an RMON Alarm

Administration > RMON

Step: 1. Configure Global Action: Add

☒ Alarm ☐ Event

Index (1-65535)

Variable

Interval (1-31622400) sec

Sample Type

Rising Threshold (0-2147483647)

Rising Event Index (0-65535)

Falling Threshold (0-2147483647)

Falling Event Index (0-65535)

Owner

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 280: Showing Configured RMON Alarms

Administration > RMON

Step: 1. Configure Global Action: Show

☒ Alarm ☐ Event

RMON Alarm List Total: 28

	Index	Status	Variable	Interval	Type	Last Value	Rising Threshold	Rising Event Index	Falling Threshold	Falling Event Index	Owner
☐	1	Valid	1.3.6.1.2.1.16.1.1.1.6.1	30	Delta	0	892800	0	446400	0	
☐	2	Valid	1.3.6.1.2.1.16.1.1.1.6.2	30	Delta	0	892800	0	446400	0	
☐	3	Valid	1.3.6.1.2.1.16.1.1.1.6.3	30	Delta	0	892800	0	446400	0	
☐	4	Valid	1.3.6.1.2.1.16.1.1.1.6.4	30	Delta	0	892800	0	446400	0	
☐	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 384](#)).
 - **Trap** – Sends a trap message to all configured trap managers (see [“Specifying Trap Managers” on page 432](#)).
 - **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 425](#)) prior to configuring it here. (Range: 1-127 characters)
- ◆ **Description** – A comment that describes this event. (Range: 1-127 characters)
- ◆ **Owner** – Name of the person who created this entry. (Range: 1-127 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 281: Configuring an RMON Event

The screenshot shows a web interface for configuring an RMON event. The breadcrumb navigation at the top reads "Administration > RMON". Below this, there are two dropdown menus: "Step:" with "1. Configure Global" selected, and "Action:" with "Add" selected. The main configuration area has two radio buttons: "Alarm" (unselected) and "Event" (selected). Below the radio buttons are several input fields: "Index (1-65535)" with the value "2", "Type" with a dropdown menu showing "Log and Trap", "Community" with the value "private", "Description" with the value "for software group", and "Owner" with the value "david". At the bottom right of the form 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 282: Showing Configured RMON Events

Index	Status	Type	Community	Description	Owner	Last Fired
1	Valid	None		None	None	00:00:00
2	Valid	Log		Log	Log	00:00:00
3	Valid	Trap	Trap	Trap	Trap	00:00:00
4	Valid	Log and Trap	Log and Trap	Log and Trap	Log and Trap	00:00:00

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 126](#).
- ◆ 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: 50)
The number of buckets granted are displayed on the Show page.
- ◆ **Owner** - Name of the person who created this entry. (Range: 1-127 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 283: Configuring an RMON History Sample

The screenshot shows the 'Administration > RMON' web interface. At the top, there is a breadcrumb trail 'Administration > RMON'. Below it, there are two dropdown menus: 'Step: 2. Configure Interface' and 'Action: Add'. Underneath these, there are two radio buttons: 'History' (selected) and 'Statistics'. Below the radio buttons, there is a 'Port' dropdown menu with the value '2'. Below the port, there are four input fields: 'Index (1-65535)' with the value '100', 'Interval (1-3600)' with the value '60' and a 'sec' unit label, 'Buckets (1-65535)' with the value '10', and 'Owner' with the value 'david'. At the bottom right, there are two buttons: 'Apply' and '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 284: Showing Configured RMON History Samples

Administration > RMON

Step: 2. Configure Interface Action: Show

History Statistics

Port 1

RMON History Port List Total: 2

Index	Status	Interval	Requested Buckets	Granted Buckets	Owner
1	Valid	1800	8	8	
2	Valid	30	8	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 285: Showing Collected RMON History Samples

Administration > RMON

Step: 2. Configure Interface Action: Show Details

History Statistics

Port 1

RMON History Details Port List Total: 9

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	00:00:01	756105	3218	91	894	0	0	0	0	0	0	0	0
2	71	00:35:01	21490	76	0	15	0	0	0	0	0	0	0	0
2	72	00:35:31	46521	120	0	15	0	0	0	0	0	0	0	0
2	73	00:36:01	21682	79	0	15	0	0	0	0	0	0	0	0
2	74	00:36:31	21554	77	0	15	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: 1-127 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 286: Configuring an RMON Statistical Sample

Administration > RMON

Step: 2. Configure Interface Action: Add

☐ History ☒ Statistics

Port 2

Index (1-65535) 100

Owner mary

Apply Revert

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 287: Showing Configured RMON Statistical Samples

Administration > RMON

Step: 2. Configure Interface Action: Show

☐ History ☒ Statistics

Port 2

RMON Statistics Port List Total: 2


	Index	Status	Owner
☐	1	Valid	abc
☐	2	Valid	test

Delete Revert

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 288: Showing Collected RMON Statistical Samples



Switch Clustering

Switch clustering is a method of grouping switches together to enable centralized management through a single unit. Switches that support clustering can be grouped together regardless of physical location or switch type, as long as they are connected to the same local network.

Command Usage

- ◆ A switch cluster has a primary unit called the "Commander" which is used to manage all other "Member" switches in the cluster. The management station can use either Telnet or the web interface to communicate directly with the Commander through its IP address, and then use the Commander to manage Member switches through the cluster's "internal" IP addresses.
- ◆ Clustered switches must be in the same Ethernet broadcast domain. In other words, clustering only functions for switches which can pass information between the Commander and potential Candidates or active Members through VLAN 4093.
- ◆ Once a switch has been configured to be a cluster Commander, it automatically discovers other cluster-enabled switches in the network. These "Candidate" switches only become cluster Members when manually selected by the administrator through the management station.
- ◆ There can be up to 100 candidates and 36 member switches in one cluster.
- ◆ A switch can only be a member of one cluster.

- ◆ The cluster VLAN 4093 is not configured by default. Before using clustering, take the following actions to set up this VLAN:
 1. Create VLAN 4093 (see [“Configuring VLAN Groups” on page 162](#)).
 2. Add the participating ports to this VLAN (see [“Adding Static Members to VLANs” on page 165](#)), and set them to hybrid mode, tagged members, PVID = 1, and acceptable frame type = all.
- ◆ After the Commander and Members have been configured, any switch in the cluster can be managed from the web agent by choosing the desired Member ID from the Show Member page.

Configuring General Settings for Clusters

Use the Administration > Cluster (Configure Global) page to create a switch cluster.

Command Usage

First be sure that clustering is enabled on the switch (the default is disabled), then set the switch as a Cluster Commander. Set a Cluster IP Pool that does not conflict with the network IP subnet. Cluster IP addresses are assigned to switches when they become Members and are used for communication between Member switches and the Commander.

Parameters

These parameters are displayed:

- ◆ **Cluster Status** – Enables or disables clustering on the switch. (Default: Disabled)
- ◆ **Commander Status** – Enables or disables the switch as a cluster Commander. (Default: Disabled)
- ◆ **IP Pool** – An “internal” IP address pool that is used to assign IP addresses to Member switches in the cluster. Internal cluster IP addresses are in the form 10.x.x.member-ID. Only the base IP address of the pool needs to be set since Member IDs can only be between 1 and 36. Note that you cannot change the cluster IP pool when the switch is currently in Commander mode. Commander mode must first be disabled. (Default: 10.254.254.1)
- ◆ **Role** – Indicates the current role of the switch in the cluster; either Commander, Member, or Candidate. (Default: Candidate)
- ◆ **Number of Members** – The current number of Member switches in the cluster.
- ◆ **Number of Candidates** – The current number of Candidate switches discovered in the network that are available to become Members.

Web Interface

To configure a switch cluster:

1. Click Administration, Cluster.
2. Select Configure Global from the Step list.
3. Set the required attributes for a Commander or a managed candidate.
4. Click Apply

Figure 289: Configuring a Switch Cluster

Administration > Cluster

Step: 1. Configure Global

Cluster Status	☒ Enabled
Commander Status	☒ Enabled
IP Pool	<input type="text" value="10.254.254.1"/>
Role	Commander
Number of Members	2
Number of Candidates	3

Apply Revert

Cluster Member Configuration Use the Administration > Cluster (Configure Member - Add) page to add Candidate switches to the cluster as Members.

Parameters

These parameters are displayed:

- ◆ **Member ID** – Specify a Member ID number for the selected Candidate switch. (Range: 1-36)
- ◆ **MAC Address** – Select a discovered switch MAC address from the Candidate Table, or enter a specific MAC address of a known switch.

Web Interface

To configure cluster members:

1. Click Administration, Cluster.
2. Select Configure Member from the Step list.
3. Select Add from the Action list.
4. Select one of the cluster candidates discovered by this switch, or enter the MAC address of a candidate.
5. Click Apply.

Figure 290: Configuring a Cluster Members

Administration > Cluster

Step: 2. Configure Member Action: Add

Member ID (1-36)

MAC Address ☒ Candidate 11-22-33-44-55-11

☐

Apply Revert

To show the cluster members:

1. Click Administration, Cluster.
2. Select Configure Member from the Step list.
3. Select Show from the Action list.

Figure 291: Showing Cluster Members

Administration > Cluster

Step: 2. Configure Member Action: Show

Cluster Member List Total: 2


	Member ID	Role	IP Address	MAC Address	Description
☐	1	Active Member	10.254.254.2	11-22-33-44-55-33	GTL-2891
☐	2	Candidate	10.254.254.3	11-22-33-44-55-77	GTL-2891

Delete Revert

To show cluster candidates:

1. Click Administration, Cluster.
2. Select Configure Member from the Step list.
3. Select Show Candidate from the Action list.

Figure 292: Showing Cluster Candidates



Role	MAC Address	Description
Candidate	11-22-33-44-55-11	GTL-2891
Active Member	11-22-33-44-55-22	GTL-2891
Candidate	11-22-33-44-55-33	GTL-2891
Candidate	11-22-33-44-55-44	GTL-2891

Managing Cluster Members

Use the Administration > Cluster (Show Member) page to manage another switch in the cluster.

Parameters

These parameters are displayed:


- ◆ **Member ID** – The ID number of the Member switch. (Range: 1-36)
- ◆ **Role** – Indicates the current status of the switch in the cluster.
- ◆ **IP Address** – The internal cluster IP address assigned to the Member switch.
- ◆ **MAC Address** – The MAC address of the Member switch.
- ◆ **Description** – The system description string of the Member switch.
- ◆ **Operate** – Remotely manage a cluster member.

Web Interface

To manage a cluster member:

1. Click Administration, Cluster.
2. Select Show Member from the Step list.
3. Select an entry from the Cluster Member List.
4. Click Operate.

Figure 293: Managing a Cluster Member



The screenshot shows a web interface for managing a cluster. At the top, it says 'Administration > Cluster'. Below that, a 'Step:' dropdown is set to '3. Show Member'. The main section is titled 'Cluster Member List' with a 'Total: 2' indicator. It contains a table with 6 columns: Member ID, Role, IP Address, MAC Address, and Description. There are two rows: Member 1 is an 'Active Member' with IP 10.254.254.2 and MAC 11-22-33-44-55-33; Member 2 is a 'Candidate' with IP 10.254.254.3 and MAC 11-22-33-44-55-77. Both have the description 'GTL-2891'. An 'Operate' button is at the bottom right.

Member ID	Role	IP Address	MAC Address	Description
1	Active Member	10.254.254.2	11-22-33-44-55-33	GTL-2891
2	Candidate	10.254.254.3	11-22-33-44-55-77	GTL-2891

Ethernet Ring Protection Switching



Note: Information in this section is based on ITU-T G.8032/Y.1344.

The ITU G.8032 recommendation specifies a protection switching mechanism and protocol for Ethernet layer network rings. Ethernet rings can provide wide-area multipoint connectivity more economically due to their reduced number of links. The mechanisms and protocol defined in G.8032 achieve highly reliable and stable protection; and never form loops, which would fatally affect network operation and service availability.

The G.8032 recommendation, also referred to as Ethernet Ring Protection Switching (ERPS), can be used to increase the availability and robustness of Ethernet rings. An Ethernet ring built using ERPS can provide resilience at a lower cost and than that provided by SONET or EAPS rings.

ERPS is more economical than EAPS in that only one physical link is required between each node in the ring. However, since it can tolerate only one break in the ring, it is not as robust as EAPS. ERPS supports up to 255 nodes in the ring structure. ERPS requires a higher convergence time when more than 16 nodes are used, but should always run under than 500 ms.

Operational Concept

Loop avoidance in the ring is achieved by guaranteeing that, at any time, traffic may flow on all but one of the ring links. This particular link is called the ring protection link (RPL), and under normal conditions this link is blocked to traffic. One designated node, the RPL owner, is responsible for blocking traffic over the RPL. When a ring failure occurs, the RPL owner is responsible for unblocking the RPL, allowing this link to be used for traffic.

Ring nodes may be in one of two states:

Idle – normal operation, no link/node faults detected in ring

Protection – Protection switching in effect after identifying a signal fault


In Idle state, the physical topology has all nodes connected in a ring. The logical topology guarantees that all nodes are connected without a loop by blocking the RPL. Each link is monitored by its two adjacent nodes using Connectivity Fault Management (CFM) protocol messages.

Protection switching (opening the RPL to traffic) occurs when a signal failure message generated by the Connectivity Fault Management (CFM) protocol is declared on one of the ring links, and the detected failure has a higher priority than any other request; or a Ring – Automatic Protection Switching protocol request (R-APS, as defined in Y.1731) is received which has a higher priority than any other local request.

A link/node failure is detected by the nodes adjacent to the failure. These nodes block the failed link and report the failure to the ring using R-APS (SF) messages. This message triggers the RPL owner to unblock the RPL, and all nodes to flush their forwarding database. The ring is now in protection state, but it remains connected in a logical topology.

When the failed link recovers, the traffic is kept blocked on the nodes adjacent to the recovered link. The nodes adjacent to the recovered link transmit R-APS (NR - no request) message indicating they have no local request. When the RPL owner receives an R-APS (NR) message it starts the Wait-To-Recover (WTR) timer. Once WTR timer expires, the RPL owner blocks the RPL and transmits an R-APS (NR, RB - ring blocked) message. Nodes receiving this message flush the forwarding database and unblock their previously blocked ports. The ring is now returned to Idle state.

Figure 294: ERPS Ring Components



Multi-ring/Ladder Network – ERPSv2 also supports multipoint-to-multipoint connectivity within interconnected rings, called a “multi-ring/ladder network” topology. This arrangement consists of conjoined rings connected by one or more interconnection points, and is based on the following criteria:

- ◆ The R-APS channels are not shared across Ethernet Ring interconnections.
- ◆ On each ring port, each traffic channel and each R-APS channel are controlled (e.g., for blocking or flushing) by the Ethernet Ring Protection Control Process (ERP Control Process) of only one ring.

- ◆ Each Major Ring or Sub-Ring must have its own RPL.


[Figure 295 on page 458](#) (Normal Condition) depicts an example of a multi-ring/ladder network. If the network is in normal operating condition, the RPL owner node of each ring blocks the transmission and reception of traffic over the RPL for that ring. This figure presents the configuration when no failure exists on any ring link.

In the figure for the Normal Condition there are two interconnected rings. Ring ERP1 is composed of ring nodes A, B, C and D and the ring links between these nodes. Ring ERP2 is composed of ring nodes C, D, E and F and the ring links C-to-F, F-to-E, E-to-D. The ring link between D and C is used for traffic on rings ERP1 and ERP2. On their own ERP2 ring links do not form a closed loop. A closed loop may be formed by the ring links of ERP2 and the ring link between the interconnection nodes that is controlled by ERP1. ERP2 is a sub-ring. Ring node A is the RPL owner node for ERP1, and ring node E is the RPL owner node for ERP2. These ring nodes (A and E) are responsible for blocking the traffic channel on the RPL for ERP1 and ERP2 respectively. There is no restriction on which ring link on an ring may be set as the RPL. For example the RPL of ERP1 could be set as the link between ring node C and D.

Ring nodes C and D, that are common to both ERP1 and ERP2, are called interconnection nodes. The ring link between the interconnection nodes are controlled and protected by the ring it belongs to. In the example for the Normal Condition, the ring link between ring nodes C and D is part of ERP1, and, as such, are controlled and protected by ERP1. Ethernet characteristic information traffic corresponding to the traffic channel may be transferred over a common Ethernet connection for ERP1 and ERP2 through the interconnection nodes C and D. Interconnection nodes C and D have separate ERP Control Processes for each Ethernet Ring.

[Figure 295 on page 458](#) (Signal Fail Condition) illustrates a situation where protection switching has occurred due to an SF condition on the ring link between interconnection nodes C and D. The failure of this ring link triggers protection only on the ring to which it belongs, in this case ERP1. The traffic and R-APS channels are blocked bi-directionally on the ports where the failure is detected and bi-directionally unblocked at the RPL connection point on ERP1. The traffic channels remain bi-directionally blocked at the RPL connection point on ERP2. This prevents the formation of a loop.

Figure 295: Ring Interconnection Architecture (Multi-ring/Ladder Network)



Configuration Guidelines for ERPS

1. Create an ERPS ring ([Configure Domain – Add](#)): The ring name is used as an index in the G.8032 database.
2. Configure the east and west interfaces ([Configure Domain – Configure Details](#)): Each node on the ring connects to it through two ring ports. Configure one port connected to the next node in the ring to the east (or clockwise direction) and another port facing west in the ring.
3. Configure the RPL owner ([Configure Domain – Configure Details](#)): Configure one node in the ring as the Ring Protection Link (RPL) owner. When this switch is configured as the RPL owner, the west ring port is set as being connected to the RPL. Under normal operations (Idle state), the RPL is blocked to ensure that a loop cannot form in the ring. If a signal failure brings down any other link in the ring, the RPL will be unblocked (Protection state) to ensure proper connectivity among all ring nodes until the failure is recovered.
4. Configure ERPS timers ([Configure Domain – Configure Details](#)): Set the Guard timer to prevent ring nodes from receiving outdated R-APS messages, the Hold-off timer to filter out intermittent link faults, and the WTR timer to verify that the ring has stabilized before blocking the RPL after recovery from a signal failure.
5. Configure the ERPS control VLAN ([Configure Domain – Configure Details](#)): Specify the control VLAN (CVLAN) used to pass R-APS ring maintenance commands. The CVLAN must NOT be configured with an IP address. In addition, only ring ports may be added to the CVLAN (prior to configuring the VLAN as a CVLAN). No other ports can be members of this VLAN (once set as a CVLAN). Also, the ring ports of the CVLAN must be tagged. Failure to observe these restrictions can result in a loop in the network.

6. Enable ERPS ([Configure Global](#)): Before enabling a ring as described in the next step, first globally enable ERPS on the switch. If ERPS has not yet been enabled or has been disabled, no ERPS rings will work.
7. Enable an ERPS ring ([Configure Domain – Configure Details](#)): Before an ERPS ring can work, it must be enabled. When configuration is completed and the ring enabled, R-APS messages will start flowing in the control VLAN, and normal traffic will begin to flow in the data VLANs. A ring can be stopped by disabling the Admin Status on any node.
8. Display ERPS status information ([Configure Domain – Show](#)): Display ERPS status information for all configured rings.

Configuration Limitations for ERPS

The following configuration limitations apply to ERPS:

- ◆ The switch supports up to six ERPS rings – each ring must have one Control VLAN, and at most 255 Data VLANs.
- ◆ Ring ports can not be a member of a trunk, nor an LACP-enabled port.
- ◆ Dynamic VLANs are not supported as protected data ports.
- ◆ Exclusive use of **STP** or **ERPS** on any port.
- ◆ The switch takes about 350 ms to detect link-up on 1000Base-T copper ports, so the convergence time on this port type is more than 50 ms.
- ◆ One VLAN must be added to an ERPS domain as the CVLAN. This can be designated as any VLAN, other than the management VLAN. The CVLAN should only contain ring ports, and must not be configured with an IP address.

ERPS Global Configuration Use the Administration > ERPS (Configure Global) page to globally enable or disable ERPS on the switch.

Parameters

These parameters are displayed:

- ◆ **ERPS Status** – Enables ERPS on the switch. (Default: Disabled)
ERPS must be enabled globally on the switch before it can be enabled on an ERPS ring (by setting the Admin Status on the [Configure Domain – Configure Details](#) page).

Web Interface

To globally enable ERPS on the switch:

1. Click Administration, ERPS.
2. Select Configure Global from the Step list.
3. Mark the ERPS Status check box.
4. Click Apply.

Figure 296: Setting ERPS Global Status

The screenshot shows a web interface for configuring ERPS. At the top, it says 'Administration > ERPS'. Below that, there's a 'Step:' dropdown menu set to '1. Configure Global'. Underneath, the 'ERPS Status' is shown with a checked checkbox and the text 'Enabled'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

ERPS Ring Configuration

Use the Administration > ERPS (Configure Domain) pages to configure ERPS rings.

Command Usage

Ring Initialization

An ERPS ring containing one Control VLAN and one or more protected Data VLANs must be configured, and the global ERPS function enabled on the switch (see [“ERPS Global Configuration” on page 459](#)) before a ring can start running. Once enabled, the RPL owner node and non-owner node state machines will start, and the ring will enter the active state.

Limitations

When configuring a ring port, note that these ports cannot be part of a spanning tree, nor can they be members of a static or dynamic trunk.

Parameters

These parameters are displayed:

Add

- ◆ **Domain Name** – Name of an ERPS ring. (Range: 1-12 characters)
- ◆ **Domain ID** – ERPS ring identifier used in R-APS messages. (Range: 1-255)

Show

- ◆ **Domain Name** – Name of a configured ERPS ring.

- ◆ **ID** – ERPS ring identifier used in R-APS messages.
- ◆ **Admin Status** – Shows whether ERPS is enabled on the switch.
- ◆ **Ver** – Shows the ERPS version.
- ◆ **MEG Level** – The maintenance entity group (MEG) level providing a communication channel for ring automatic protection switching (R-APS) information.
- ◆ **Control VLAN** – Shows the Control VLAN ID.
- ◆ **Node State** – Shows the following ERPS states:
 - Init – The ERPS ring has started but has not yet determined the status of the ring.
 - Idle – If all nodes in a ring are in this state, it means that all the links in the ring are up. This state will switch to protection state if a link failure occurs.
 - Protection – If a node in this state, it means that a link failure has occurred. This state will switch to idle state if all the failed links recover.
- ◆ **Type** – Shows node type as None, RPL Owner or RPL Neighbor.
- ◆ **Revertive** – Shows if revertive or non-revertive recovery is selected.
- ◆ **W/E** – Shows information on the west and east ring port for this node.
- ◆ **West Port** – Shows the west ring port for this node.
- ◆ **East Port** – Shows the east ring port for this node.
- ◆ **Interface** – The port or trunk which is configured as a ring port.
- ◆ **Port State** – The operational state:
 - Blocking – The transmission and reception of traffic is blocked and the forwarding of R-APS messages is blocked, but the transmission of locally generated R-APS messages is allowed and the reception of all R-APS messages is allowed.
 - Forwarding – The transmission and reception of traffic is allowed; transmission, reception and forwarding of R-APS messages is allowed.
 - Unknown – The interface is not in a known state (includes the domain being disabled).
- ◆ **Local SF** – A signal fault generated on a link to the local node.
- ◆ **Local FS** – Shows if a forced switch command was issued on this interface.
- ◆ **Local MS** – Shows if a manual switch command was issued on this interface.

- ◆ **MEP** – The CFM MEP used to monitor the status on this link.
- ◆ **RPL** – Shows if this node is connected to the RPL.

Configure Details

- ◆ **Domain Name** – Name of a configured ERPS ring. (Range: 1-12 characters)
Service Instances within each ring are based on a unique maintenance association for the specific users, distinguished by the ring name, maintenance level, maintenance association's name, and assigned VLAN. Up to 6 ERPS rings can be configured on the switch.
- ◆ **Domain ID** – ERPS ring identifier used in R-APS messages. (Range: 1-255; Default: None)
R-APS information is carried in an R-APS PDUs. The last octet of the MAC address is designated as the Ring ID (01-19-A7-00-00-[Ring ID]). If use of the default MAC address is disabled for the R-APS Def MAC parameter, then the Domain ID will be used in R-APS PDUs.
- ◆ **Admin Status** – Activates the current ERPS ring. (Default: Disabled)
Before enabling a ring, the global ERPS function should be enabled see ("[ERPS Global Configuration](#)" on page 459), the east and west ring ports configured on each node, the RPL owner specified, and the control VLAN configured.
Once enabled, the RPL owner node and non-owner node state machines will start, and the ring will enter idle state if no signal failures are detected.
- ◆ **Version** – Specifies compatibility with the following ERPS versions:
 - 1 - ERPS version 1 based on ITU-T G.8032/Y.1344.
 - 2 - ERPS version 2 based on ITU-T G.8032/Y.1344 Version 2. (This is the default setting.)

In addition to the basic features provided by version 1, version 2 also supports:

- Multi-ring/ladder network support
- Revertive/Non-revertive recovery
- Forced Switch (FS) and Manual Switch (MS) commands for manually blocking a particular ring port
- Flush FDB (forwarding database) logic which reduces amount of flush FDB operations in the ring
- Support of multiple ERP instances on a single ring

Version 2 is backward compatible with Version 1. If version 2 is specified, the inputs and commands are forwarded transparently. If set to version 1, MS and FS operator commands are filtered, and the switch set to revertive mode.

The version number is automatically set to "1" when a ring node, supporting only the functionalities of G.8032v1, exists on the same ring with other nodes that support G.8032v2.

When ring nodes running G.8032v1 and G.8032v2 co-exist on a ring, the ring ID of each node is configured as "1".

In version 1, the MAC address 01-19-A7-00-00-01 is used for the node identifier. The R-APS Def MAC parameter has no effect.

- ◆ **MEG Level** – The maintenance entity group (MEG) level which provides a communication channel for ring automatic protection switching (R-APS) information. (Range: 0-7)

This parameter is used to ensure that received R-APS PDUs are directed for this ring. A unique level should be configured for each local ring if there are many R-APS PDUs passing through this switch.

- ◆ **Control VLAN** – A dedicated VLAN used for sending and receiving E-APS protocol messages. (Range: 1-4094)

Configure one control VLAN for each ERPS ring. First create the VLAN to be used as the control VLAN (see ["Configuring VLAN Groups" on page 162](#)), add the ring ports for the east and west interface as tagged members to this VLAN (see ["Adding Static Members to VLANs" on page 165](#)), and then use this parameter to add it to the ring.

The following restrictions are recommended to avoid creating a loop in the network or other problems which may occur under some situations:

- The Control VLAN must not be configured as a Layer 3 interface (with an IP address), a dynamic VLAN (with GVRP enabled), nor as a private VLAN.
- In addition, only ring ports may be added to the Control VLAN. No other ports can be members of this VLAN.
- Also, the ring ports of the Control VLAN must be tagged.

Once the ring has been activated, the configuration of the control VLAN cannot be modified. Use the Admin Status parameter to stop the ERPS ring before making any configuration changes to the control VLAN.

- ◆ **Node State** – Refer to the parameters for the Show page.

- ◆ **Node Type** – Shows ERPS node type as one of the following:

- **None** – Node is neither Ring Protection Link (RPL) owner nor neighbor. (This is the default setting.)
- **RPL Owner** – Specifies a ring node to be the RPL owner.
 - Only one RPL owner can be configured on a ring. The owner blocks traffic on the RPL during Idle state, and unblocks it during Protection state (that is, when a signal fault is detected on the ring or the

protection state is enabled with the Forced Switch or Manual Switch commands on the Configure Operation page).

- The east and west connections to the ring must be specified for all ring nodes. When this switch is configured as the RPL owner, the west ring port is automatically set as being connected to the RPL.
- **RPL Neighbor** – Specifies a ring node to be the RPL neighbor.
 - The RPL neighbor node, when configured, is a ring node adjacent to the RPL that is responsible for blocking its end of the RPL under normal conditions (i.e., the ring is established and no requests are present in the ring) in addition to the block at the other end by the RPL Owner Node. The RPL neighbor node may participate in blocking or unblocking its end of the RPL, but is not responsible for activating the reversion behavior.
 - Only one RPL owner can be configured on a ring. If the switch is set as the RPL owner for an ERPS domain, the west ring port is set as one end of the RPL. If the switch is set as the RPL neighbor for an ERPS domain, the east ring port is set as the other end of the RPL.
 - The east and west connections to the ring must be specified for all ring nodes. When this switch is configured as the RPL neighbor, the east ring port is set as being connected to the RPL.
 - Note that is not mandatory to declare a RPL neighbor.
- ◆ **Revertive** – Sets the method of recovery to Idle State through revertive or non-revertive mode. (Default: Enabled)
 - Revertive behavior allows the switch to automatically return the RPL from Protection state to Idle state through the exchange of protocol messages.
Non-revertive behavior for Protection, Forced Switch (FS), and Manual Switch (MS) states are basically the same. Non-revertive behavior requires the RPL to be restored from Protection state to Idle state using the Clear command (Configure Operation page).
 - Recovery for Protection Switching – A ring node that has one or more ring ports in an SF (Signal Fail) condition, upon detecting the SF condition cleared, keeps at least one of its ring ports blocked for the traffic channel and for the R-APS channel, until the RPL is blocked as a result of ring protection reversion, or until there is another higher priority request (e.g., an SF condition) in the ring.
A ring node that has one ring port in an SF condition and detects the SF condition cleared, continuously transmits the R-APS (NR – no request) message with its own Node ID as the priority information over both ring ports, informing that no request is present at this ring node and initiates a guard timer. When another recovered ring node (or nodes) holding the link block receives this message, it compares the Node ID information with its own Node ID. If the received R-APS (NR) message has the higher priority,

this ring node unblocks its ring ports. Otherwise, the block remains unchanged. As a result, there is only one link with one end blocked.

The ring nodes stop transmitting R-APS (NR) messages when they accept an R-APS (NR, RB – RPL Blocked), or when another higher priority request is received.

- **Recovery with Revertive Mode** – When all ring links and ring nodes have recovered and no external requests are active, reversion is handled in the following way:
 - a.** The reception of an R-APS (NR) message causes the RPL Owner Node to start the WTR (Wait-to-Restore) timer.
 - b.** The WTR timer is cancelled if during the WTR period a higher priority request than NR is accepted by the RPL Owner Node or is declared locally at the RPL Owner Node.
 - c.** When the WTR timer expires, without the presence of any other higher priority request, the RPL Owner Node initiates reversion by blocking its traffic channel over the RPL, transmitting an R-APS (NR, RB) message over both ring ports, informing the ring that the RPL is blocked, and performing a flush FDB action.
 - d.** The acceptance of the R-APS (NR, RB) message causes all ring nodes to unblock any blocked non-RPL link that does not have an SF condition. If it is an R-APS (NR, RB) message without a DNF (do not flush) indication, all ring nodes flush the FDB.
- **Recovery with Non-revertive Mode** – In non-revertive operation, the ring does not automatically revert when all ring links and ring nodes have recovered and no external requests are active. Non-revertive operation is handled in the following way:
 - a.** The RPL Owner Node does not generate a response on reception of an R-APS (NR) messages.
 - b.** When other healthy ring nodes receive the NR (Node ID) message, no action is taken in response to the message.
 - c.** When the operator issues the Clear command (Configure Operation page) for non-revertive mode at the RPL Owner Node, the non-revertive operation is cleared, the RPL Owner Node blocks its RPL port, and transmits an R-APS (NR, RB) message in both directions, repeatedly.
 - d.** Upon receiving an R-APS (NR, RB) message, any blocking node should unblock its non-failed ring port. If it is an R-APS (NR, RB) message without a DNF indication, all ring nodes flush the FDB.
- **Recovery for Forced Switching** – A Forced Switch command is removed by issuing the Clear command (Configure Operation page) to the same ring node where Forced Switch mode is in effect. The clear command

removes any existing local operator commands, and triggers reversion if the ring is in revertive behavior mode.

The ring node where the Forced Switch was cleared continuously transmits the R-APS (NR) message on both ring ports, informing other nodes that no request is present at this ring node. The ring nodes stop transmitting R-APS (NR) messages when they accept an RAPS (NR, RB) message, or when another higher priority request is received.

If the ring node where the Forced Switch was cleared receives an R-APS (NR) message with a Node ID higher than its own Node ID, it unblocks any ring port which does not have an SF condition and stops transmitting R-APS (NR) message over both ring ports.

- Recovery with revertive mode is handled as follows:
 - a. The reception of an R-APS (NR) message causes the RPL Owner Node to start the WTB timer.
 - b. The WTB timer is cancelled if during the WTB period a higher priority request than NR is accepted by the RPL Owner Node or is declared locally at the RPL Owner Node.
 - c. When the WTB timer expires, in the absence of any other higher priority request, the RPL Owner Node initiates reversion by blocking the traffic channel over the RPL, transmitting an R-APS (NR, RB) message over both ring ports, informing the ring that the RPL is blocked, and flushes the FDB.
 - d. The acceptance of the R-APS (NR, RB) message causes all ring nodes to unblock any blocked non-RPL that does not have an SF condition. If it is an R-APS (NR, RB) message without a DNF indication, all ring nodes flush their FDB. This action unblocks the ring port which was blocked as a result of an operator command.
- Recovery with non-revertive mode is handled as follows:
 - a. The RPL Owner Node, upon reception of an R-APS(NR) message and in the absence of any other higher priority request does not perform any action.
 - b. Then, after the operator issues the Clear command (Configure Operation page) at the RPL Owner Node, this ring node blocks the ring port attached to the RPL, transmits an R-APS (NR, RB) message on both ring ports, informing the ring that the RPL is blocked, and flushes its FDB.
 - c. The acceptance of the R-APS (NR, RB) message triggers all ring nodes to unblock any blocked non-RPL which does not have an SF condition. If it is an R-APS (NR, RB) message without a DNF indication, all ring nodes flush their FDB. This action unblocks

the ring port which was blocked as result of an operator command.

- Recovery for Manual Switching – A Manual Switch command is removed by issuing the Clear command (Configure Operation page) at the same ring node where the Manual Switch is in effect. The clear command removes any existing local operator commands, and triggers reversion if the ring is in revertive behavior mode.

The ring node where the Manual Switch was cleared keeps the ring port blocked for the traffic channel and for the R-APS channel, due to the previous Manual Switch command. This ring port is kept blocked until the RPL is blocked as a result of ring protection reversion, or until there is another higher priority request (e.g., an SF condition) in the ring.

The Ethernet Ring Node where the Manual Switch was cleared continuously transmits the R-APS (NR) message on both ring ports, informing that no request is present at this ring node. The ring nodes stop transmitting R-APS (NR) messages when they accept an RAPS (NR, RB) message, or when another higher priority request is received.

If the ring node where the Manual Switch was cleared receives an R-APS (NR) message with a Node ID higher than its own Node ID, it unblocks any ring port which does not have an SF condition and stops transmitting R-APS (NR) message on both ring ports.

- Recovery with revertive mode is handled as follows:
 - a. The RPL Owner Node, upon reception of an R-APS (NR) message and in the absence of any other higher priority request, starts the WTB timer and waits for it to expire. While the WTB timer is running, any latent R-APS (MS) message is ignored due to the higher priority of the WTB running signal.
 - b. When the WTB timer expires, it generates the WTB expire signal. The RPL Owner Node, upon reception of this signal, initiates reversion by blocking the traffic channel on the RPL, transmitting an R-APS (NR, RB) message over both ring ports, informing the ring that the RPL is blocked, and flushes its FDB.
 - c. The acceptance of the R-APS (NR, RB) message causes all ring nodes to unblock any blocked non-RPL that does not have an SF condition. If it is an R-APS (NR, RB) message without a DNF indication, all Ethernet Ring Nodes flush their FDB. This action unblocks the ring port which was blocked as a result of an operator command.

- Recovery with non-revertive mode is handled as follows:
 - a. The RPL Owner Node, upon reception of an R-APS (NR) message and in the absence of any other higher priority request does not perform any action.
 - b. Then, after the operator issues the Clear command (Configure Operation page) at the RPL Owner Node, this ring node blocks the ring port attached to the RPL, transmits an R-APS (NR, RB) message over both ring ports, informing the ring that the RPL is blocked, and flushes its FDB.
 - c. The acceptance of the R-APS (NR, RB) message triggers all ring nodes to unblock any blocked non-RPL which does not have an SF condition. If it is an R-APS (NR, RB) message without a DNF indication, all ring nodes flush their FDB. This action unblocks the ring port which was blocked as result of an operator command.

◆ **Major Domain** – The ERPS ring used for sending control packets.

This switch can support up to six rings. However, ERPS control packets can only be sent on one ring. This parameter is used to indicate that the current ring is a secondary ring, and to specify the major ring which will be used to send ERPS control packets.

The Ring Protection Link (RPL) is always the west port. So the physical port on a secondary ring must be the west port. In other words, if a domain has two physical ring ports, this ring can only be a major ring, not a secondary ring (or sub-domain) which can have only one physical ring port. The major domain therefore cannot be set if the east port is already configured.

◆ **Node ID** – A MAC address unique to the ring node. The MAC address must be specified in the format xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx. (Default: CPU MAC address)

The ring node identifier is used to identify a node in R-APS messages for both automatic and manual switching recovery operations.

For example, a node that has one ring port in SF condition and detects that the condition has been cleared, will continuously transmit R-APS (NR) messages with its own Node ID as priority information over both ring ports, informing its neighbors that no request is present at this node. When another recovered node holding the link blocked receives this message, it compares the Node ID information with its own. If the received R-APS (NR) message has a higher priority, this unblocks its ring ports. Otherwise, the block remains unchanged.


The node identifier may also be used for debugging, such as to distinguish messages when a node is connected to more than one ring.

◆ **R-APS with VC** – Configures an R-APS virtual channel to connect two interconnection points on a sub-ring, allowing ERPS protocol traffic to be tunneled across an arbitrary Ethernet network. (Default: Enabled)

- A sub-ring may be attached to a primary ring with or without a virtual channel. A virtual channel is used to connect two interconnection points on the sub-ring, tunneling R-APS control messages across an arbitrary Ethernet network topology. If a virtual channel is not used to cross the intermediate Ethernet network, data in the traffic channel will still flow across the network, but the all R-APS messages will be terminated at the interconnection points.
- Sub-ring with R-APS Virtual Channel – When using a virtual channel to tunnel R-APS messages between interconnection points on a sub-ring, the R-APS virtual channel may or may not follow the same path as the traffic channel over the network. R-APS messages that are forwarded over the sub-ring's virtual channel are broadcast or multicast over the interconnected network. For this reason the broadcast/multicast domain of the virtual channel should be limited to the necessary links and nodes. For example, the virtual channel could span only the interconnecting rings or sub-rings that are necessary for forwarding R-APS messages of this sub-ring. Care must also be taken to ensure that the local RAPS messages of the sub-ring being transported over the virtual channel into the interconnected network can be uniquely distinguished from those of other interconnected ring R-APS messages. This can be achieved by, for example, by using separate VIDs for the virtual channels of different sub-rings.

Note that the R-APS virtual channel requires a certain amount of bandwidth to forward R-APS messages on the interconnected Ethernet network where a sub-ring is attached. Also note that the protection switching time of the sub-ring may be affected if R-APS messages traverse a long distance over an R-APS virtual channel.


Figure 297: Sub-ring with Virtual Channel



- Sub-ring without R-APS Virtual Channel – Under certain circumstances it may not be desirable to use a virtual channel to interconnect the sub-ring over an arbitrary Ethernet network. In this situation, the R-APS messages are terminated on the interconnection points. Since the sub-ring does not provide an R-APS channel nor R-APS virtual channel beyond the interconnection points, R-APS channel blocking is not employed on the normal ring links to avoid channel segmentation. As a result, a failure at any ring link in the sub-ring will cause the R-APS channel of the sub-ring to be segmented, thus preventing R-APS message exchange between some of the sub-ring's ring nodes.

No R-APS messages are inserted or extracted by other rings or sub-rings at the interconnection nodes where a sub-ring is attached. Hence there is no need for either additional bandwidth or for different VIDs/Ring IDs for the ring interconnection. Furthermore, protection switching time for a sub-ring is independent from the configuration or topology of the interconnected rings. In addition, this option always ensures that an interconnected network forms a tree topology regardless of its interconnection configuration. This means that it is not necessary to take precautions against forming a loop which is potentially composed of a whole interconnected network.

Figure 298: Sub-ring without Virtual Channel



- ◆ **R-APS Def MAC** – Sets the switch’s MAC address to be used as the node identifier in R-APS messages. (Default: Enabled)

When ring nodes running ERPSv1 and ERPSv2 co-exist on the same ring, the Ring ID of each ring node must be configured as “1”.

If this command is disabled, the following strings are used as the node identifier:

- ERPSv1: 01-19-A7-00-00-01
- ERPSv2: 01-19-A7-00-00-[Ring ID]

- ◆ **Propagate TC** – Enables propagation of topology change messages from a secondary ring to the primary ring. (Default: Disabled)

When a secondary ring detects a topology change, it can pass a message about this event to the major ring. When the major ring receives this kind of message from a secondary ring, it can clear the MAC addresses on its ring ports to help the secondary ring restore its connections more quickly through protection switching.


When the MAC addresses are cleared, data traffic may flood onto the major ring. The data traffic will become stable after the MAC addresses are learned again. The major ring will not be broken, but the bandwidth of data traffic on the major ring may suffer for a short period of time due to this flooding behavior.

- ◆ **Non-ERPS Device Protection** – Sends non-standard health-check packets when an owner node enters protection state without any link down event having been detected through Signal Fault messages. (Default: Disabled)

- The RPL owner node detects a failed link when it receives R-APS (SF - signal fault) messages from nodes adjacent to the failed link. The owner then enters protection state by unblocking the RPL. However, using this standard recovery procedure may cause a non-ERPS device to become isolated when the ERPS device adjacent to it detects a continuity check message (CCM) loss event and blocks the link between the non-ERPS device and ERPS device.

CCMs are propagated by the Connectivity Fault Management (CFM) protocol as described under [“Connectivity Fault Management” on page 480](#). If the standard recovery procedure were used as shown in the following figure, and node E detected CCM loss, it would send an R-APS (SF) message to the RPL owner and block the link to node D, isolating that non-ERPS device.

Figure 299: Non-ERPS Device Protection



When non-ERPS device protection is enabled on the ring, the ring ports on the RPL owner node and non-owner nodes will not be blocked when signal loss is detected by CCM loss events.

- When non-ERPS device protection is enabled on an RPL owner node, it will send non-standard health-check packets to poll the ring health when it enters the protection state. It does not use the normal procedure of waiting to receive an R-APS (NR - no request) message from nodes adjacent to the recovered link. Instead, it waits to see if the non-standard health-check packets loop back. If they do, indicating that the fault has been resolved, the RPL will be blocked.

After blocking the RPL, the owner node will still transmit an R-APS (NR, RB - ring blocked) message. ERPS-compliant nodes receiving this message flush their forwarding database and unblock previously blocked ports. The ring is now returned to Idle state.

- ◆ **Holdoff Timer** – The hold-off timer is used to filter out intermittent link faults. Faults will only be reported to the ring protection mechanism if this timer expires. (Range: 0-10000 milliseconds, in steps of 100 milliseconds)

In order to coordinate timing of protection switches at multiple layers, a hold-off timer may be required. Its purpose is to allow, for example, a server layer protection switch to have a chance to fix the problem before switching at a client layer.

When a new defect or more severe defect occurs (new Signal Failure), this event will not be reported immediately to the protection switching mechanism if the provisioned hold-off timer value is non-zero. Instead, the hold-off timer will be started. When the timer expires, whether a defect still exists or not, the timer

will be checked. If one does exist, that defect will be reported to the protection switching mechanism. The reported defect need not be the same one that started the timer.

- ◆ **Guard Timer** – The guard timer is used to prevent ring nodes from receiving outdated R-APS messages. During the duration of the guard timer, all received R-APS messages are ignored by the ring protection control process, giving time for old messages still circulating on the ring to expire. (Range: 10-2000 milliseconds, in steps of 10 milliseconds)

The guard timer duration should be greater than the maximum expected forwarding delay for an R-APS message to pass around the ring. A side-effect of the guard timer is that during its duration, a node will be unaware of new or existing ring requests transmitted from other nodes.

- ◆ **WTB Timer** – The Wait to Block (WTB) timer is used when clearing Forced Switch (FS) and Manual Switch (MS) commands. As multiple FS commands are allowed to co-exist in a ring, the WTB timer ensures that clearing of a single FS command does not trigger re-blocking of the RPL. When clearing an MS command, the WTB timer prevents the formation of a closed loop due to possible a timing anomaly where the RPL owner node receives an outdated remote MS request during the recovery process.

When recovering from an FS or MS command, the delay timer must be long enough to receive any latent remote FS or MS commands. This delay timer called the WTB timer is defined to be 5 seconds longer than the guard timer. This is enough time to allow a reporting ring node to transmit two R-APS messages and allow the ring to identify the latent condition.

This delay timer is activated on the RPL owner node. When the relevant delay timer expires, the RPL owner node initiates the reversion process by transmitting an R-APS (NR, RB) message. The delay timer, (i.e., WTR or WTB) is deactivated when any higher priority request preempts this delay timer.

The delay timers (i.e. WTR and WTB) may be started and stopped by the system. A request to start running the delay timer does not restart the delay timer. A request to stop the delay timer stops the delay timer and resets its value. The Clear command (Configure Operation page) can be used to stop the delay timer.

- ◆ **WTR Timer** – The wait-to-restore timer is used to verify that the ring has stabilized before blocking the RPL after recovery from a signal failure. (Range: 5-12 minutes)

If the switch goes into ring protection state due to a signal failure, after the failure condition is cleared, the RPL owner will start the wait-to-restore timer and wait until it expires to verify that the ring has stabilized before blocking the RPL and returning to the Idle (normal operating) state.

- ◆ **WTB Expire** – The time before the wait-to-block timer expires.

- ◆ **WTR Expire** – The time before the wait-to-restore timer expires.

- ◆ **West/East** – Connects to next ring node to the west/east.

Each node must be connected to two neighbors on the ring. For convenience, the ports connected are referred to as east and west ports. Alternatively, the closest neighbor to the east should be the next node in the ring in a clockwise direction, and the closest neighbor to the west should be the next node in the ring in a counter-clockwise direction.

- ◆ **Interface** – The port or trunk attached to the west or east ring port.

Note that a ring port cannot be configured as a member of a spanning tree, a dynamic trunk, or a static trunk.

- ◆ **Port State** – Once configured, this field shows the operational state of the ring ports for this node:

- **Blocking** – The transmission and reception of traffic is blocked and the forwarding of R-APS messages is blocked, but the transmission of locally generated R-APS messages is allowed and the reception of all R-APS messages is allowed.
- **Forwarding** – The transmission and reception of traffic is allowed; transmission, reception and forwarding of R-APS messages is allowed.
- **Unknown** – The interface is not in a known state.

- ◆ **Local SF** – Shows if a signal fault exists on a link to the local node.

- ◆ **Local FS** – Shows if a forced switch command was issued on this interface.

- ◆ **Local MS** – Shows if a manual switch command was issued on this interface.

- ◆ **MEP** – Specifies the CCM MEPs used to monitor the link on a ring node.

If a MEP is used to monitor the link status of an ERPS node with CFM continuity check messages, then the MEG Level parameter on this configuration page must match the authorized maintenance level of the CFM domain to which the specified MEP belongs. (See [“Configuring CFM Maintenance Domains” on page 487.](#))

To ensure complete monitoring of a ring node, specify the CFM MEPs used to monitor both the east and west ports of the ring node.

If CFM determines that a MEP node which has been configured to monitor a ring port with this command has gone down, this information is passed to ERPS, which in turn processes it as a ring node failure. For more information on how ERPS recovers from a node failure, refer to the description of the Revertive parameter on this configuration page.

- ◆ **RPL** – If node is connected to the RPL, this shows by which interface.

Web Interface

To create an ERPS ring:

1. Click Administration, ERPS.
2. Select Configure Domain from the Step list.
3. Select Add from the Action list.
4. Enter a name and optional identifier for the ring.
5. Click Apply.

Figure 300: Creating an ERPS Ring

Administration > ERPS

Step: 2. Configure Domain Action: Add

Domain Name

Domain ID (1-255) ☒

Apply Revert

To configure the ERPS parameters for a ring:

1. Click Administration, ERPS.
2. Select Configure Domain from the Step list.
3. Select Configure Details from the Action list.
4. Configure the ERPS parameters for this node. Note that spanning tree protocol cannot be configured on the ring ports, nor can these ports be members of a static or dynamic trunk. And the control VLAN must be unique for each ring. Adjust the protocol timers as required. The RPL owner must be set on one of the rings. And the administrative status enabled once all of the other settings have been entered.
5. Click Apply.

Figure 301: Creating an ERPS Ring

Administration > ERPS

Step: 2. Configure Domain Action: Configure Details

Domain Name: rd1

Domain ID: 1

Admin Status: ☒ Enabled

Version: 2

MEG Level (0-7): 1

Control VLAN: ☒ 2

Node State: Idle

Node Type: RPL Owner

Revertive: ☒ Enabled

Major Domain: ☐

Node ID: 00-E0-0C-00-00-FD

R-APS with VC: ☒ Enabled

R-APS Def MAC: ☒ Enabled

Propagate TC: ☐ Enabled

Non-ERPS Dev Protect: ☐ Enabled

Holdoff Timer (0-10000): 0 ms

Guard Timer (10-2000): 500 ms

WTB Timer: 5500 ms

WTR Timer (5-12): 5 min

WTB Expire:

WTR Expire:

West: ☒ Enabled

East: ☒ Enabled

Interface: Eth 1/10

Interface: Eth 1/12

Port State: Blocking

Port State: Forwarding

Local SF: No

Local SF: No

Local FS: No

Local FS: No

Local MS: No

Local MS: No

MEP (1-8191): ☐

MEP (1-8191): ☐

RPL: Yes

RPL: No

Apply Revert

To show the configured ERPS rings:

1. Click Administration, ERPS.
2. Select Configure Domain from the Step list.
3. Select Show from the Action list.

Figure 302: Showing Configured ERPS Rings

Administration > ERPS

Step: 2. Configure Domain Action: Show

Domain List Total: 1

	Domain Name	ID	Admin Status	Ver	MEG Level	Control VLAN	Node State	Type	Revertive	W/E	Interface	Port State	Local SF	Local FS	Local MS	MEP	RPL
☐	rd1	1	Enabled	2	1	2	Idle	RPL Owner	Yes	West	Eth 1/10	Blocking	No	No	No		Yes
										East	Eth 1/12	Forwarding	No	No	No		No

Delete Revert

ERPS Forced and Manual Mode Operations Use the Administration > ERPS (Configure Operation) page to block a ring port using Forced Switch or Manual Switch commands.

Parameters

These parameters are displayed:

- ◆ **Domain Name** – Name of a configured ERPS ring.
- ◆ **Operation** – Specifies a Forced Switch (FS) or Manual Switch (MS) operation on the east or west ring port.
 - **Forced Switch** – Blocks specified ring port. (Options: West or East)
 - A ring with no pending request has a logical topology with the traffic channel blocked at the RPL and unblocked on all other ring links. In this situation, the FS command triggers protection switching as follows:
 - a. The ring node where an FS command was issued blocks the traffic channel and R-APS channel on the ring port to which the command was issued, and unblocks the other ring port.
 - b. The ring node where the FS command was issued transmits R-APS messages indicating FS over both ring ports. R-APS (FS) messages are continuously transmitted by this ring node while the local FS command is the ring node's highest priority command (see [Table 31 on page 477](#)). The R-APS (FS) message informs other ring nodes of the FS command and that the traffic channel is blocked on one ring port.
 - c. A ring node accepting an R-APS (FS) message, without any local higher priority requests unblocks any blocked ring port. This action subsequently unblocks the traffic channel over the RPL.
 - d. The ring node accepting an R-APS (FS) message, without any local higher priority requests stops transmission of R-APS messages.
 - e. The ring node receiving an R-APS (FS) message flushes its FDB.
 - Protection switching on a forced switch request is completed when the above actions are performed by each ring node. At this point, traffic flows around the ring are resumed. From this point on the following rules apply regarding processing of further forced switch commands:
 - While an existing forced switch request is present in a ring, any new forced switch request is accepted, except on a ring node having a prior local forced switch request. The ring nodes where further forced switch commands are issued block the traffic channel and R-APS channel on the ring port at which the forced switch was issued. The ring node where the forced switch command was issued transmits an R-APS message over both ring ports indicating FS. R-APS (FS) messages are continuously transmitted by this ring node while the local FS command is the ring node's highest priority

command. As such, two or more forced switches are allowed in the ring, which may inadvertently cause the segmentation of an ring. It is the responsibility of the operator to prevent this effect if it is undesirable.

Ring protection requests, commands and R-APS signals have the priorities as specified in the following table.

Table 31: ERPS Request/State Priority

Request / State and Status	Type	Priority
Clear	local	highest
FS	local	
R-APS (FS)	remote	
local SF*	local	
local clear SF	local	
R-APS (SF)	remote	
R-APS (MS)	remote	
MS	local	
WTR Expires	local	
WTR Running	local	
WTB Expires	local	
WTB Running	local	
R-APS (NR, RB)	remote	
R-APS (NR)	remote	lowest

* If an Ethernet Ring Node is in the Forced Switch state, local SF is ignored.

- Recovery for forced switching under revertive and non-revertive mode is described under the Revertive parameter.
- When a ring is under an FS condition, and the node at which an FS command was issued is removed or fails, the ring remains in FS state because the FS command can only be cleared at node where the FS command was issued. This results in an unrecoverable FS condition.

When performing a maintenance procedure (e.g., replacing, upgrading) on a ring node (or a ring link), it is recommended that FS commands be issued at the two adjacent ring nodes instead of directly issuing a FS command at the ring node under maintenance in order to avoid falling into the above mentioned unrecoverable situation.

- **Manual Switch** – Blocks specified ring port, in the absence of a failure or an FS command. (Options: West or East)
- A ring with no request has a logical topology with the traffic channel blocked at the RPL and unblocked on all other ring links. In this situation, the Manual Switch command triggers protection switching as follows:
 - a. If no other higher priority commands exist, the ring node, where a manual switch command was issued, blocks the traffic channel and R-APS channel on the ring port to which the command was issued, and unblocks the other ring port.
 - b. If no other higher priority commands exist, the ring node where the manual switch command was issued transmits R-APS messages over both ring ports indicating MS. R-APS (MS) message are continuously transmitted by this ring node while the local MS command is the ring node's highest priority command (see [Table 31 on page 477](#)). The R-APS (MS) message informs other ring nodes of the MS command and that the traffic channel is blocked on one ring port.
 - c. If no other higher priority commands exist and assuming the ring node was in Idle state before the manual switch command was issued, the ring node flushes its local FDB.
 - d. A ring node accepting an R-APS (MS) message, without any local higher priority requests unblocks any blocked ring port which does not have an SF condition. This action subsequently unblocks the traffic channel over the RPL.
 - e. A ring node accepting an R-APS (MS) message, without any local higher priority requests stops transmitting R-APS messages.
 - f. A ring node receiving an R-APS (MS) message flushes its FDB.
- Protection switching on a manual switch request is completed when the above actions are performed by each ring node. At this point, traffic flows around the ring are resumed. From this point on, the following rules apply regarding processing of further manual switch commands:
 - a. While an existing manual switch request is present in the ring, any new manual switch request is rejected. The request is rejected at the ring node where the new request is issued and a notification is generated to inform the operator that the new MS request was not accepted.
 - b. A ring node with a local manual switch command which receives an R-APS (MS) message with a different Node ID clears its manual switch request and starts transmitting R-APS (NR) messages. The ring node keeps the ring port blocked due to the previous manual switch command.

- c. An ring node with a local manual switch command that receives an R-APS message or a local request of higher priority than R-APS (MS) clear its manual switch request. The ring node then processes the new higher priority request.
- Recovery for manual switching under revertive and non-revertive mode is described under the Revertive parameter.
- **Clear** – Manually clears the protection state which has been invoked by a forced switch or manual switch command, and the node is operating under non-revertive mode; or before the WTR or WTB timer expires when the node is operating in revertive mode.
 - Two steps are required to make a ring operating in non-revertive mode return to Idle state from forced switch or manual switch state:
 1. Issue a Clear command to remove the forced switch command on the node where a local forced switch command is active.
 2. Issue a Clear command on the RPL owner node to trigger the reversion.
 - The Clear command will also stop the WTR and WTB delay timers and reset their values.
 - More detailed information about using this command for non-revertive mode is included under the Revertive parameter. (See the Command Usage section under [“ERPS Ring Configuration” on page 460.](#))

Web Interface

To block a ring port:

1. Click Administration, ERPS.
2. Select Configure Domain from the Step list.
3. Select Configure Operation from the Action list.
4. Select the domain name from the drop-down list.
5. Specify a Forced Switch, Manual Switch, or Clear operation.
6. Click Apply.

Figure 303: Blocking an ERPS Ring Port

Administration > ERPS

Step: 2. Configure Domain Action: Configure Operation

Domain Name: rd1

Operation: Clear

Apply Revert

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 coarse 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 304: 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 305: 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.

Loopback 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 MEP's 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 link trace message is a multicast CFM frame initiated by a MEP, and forwarded from MIP to MIP, with each MIP generating a link trace reply, up to the point at which the link trace 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 link trace message.

Link trace 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 Link Trace 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 has 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.

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"](#).

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.
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 306: Configuring Global Settings for CFM

The screenshot displays the 'Administration > CFM' web interface. At the top, a breadcrumb trail shows 'Administration > CFM'. Below this, a 'Step:' dropdown menu is set to '1. Configure Global'. The main content area is divided into two sections: 'Global Configuration' and 'SNMP Trap Configuration'. In the 'Global Configuration' section, 'CFM Status' is checked and labeled 'Enabled'. 'MEP Cross Check Start Delay (1-65535)' is set to '10' seconds. 'Link Trace Cache' is checked and labeled 'Enabled'. 'Link Trace Cache Hold Time (1-65535)' is set to '100' minutes. 'Link Trace Cache Size (1-4095)' is set to '100' entries. The 'SNMP Trap Configuration' section lists several traps, each with an 'Enabled' checkbox: 'Connectivity Check Config', 'Connectivity Check Loop', 'Connectivity Check MEP Down', 'Connectivity Check MEP Up', 'Cross Check MA Up', 'Cross Check MEP Missing', and 'Cross Check MEP Unknown'. At the bottom right, there are 'Apply' and 'Revert' buttons.

Global Configuration	
CFM Status	☒ Enabled
MEP Cross Check Start Delay (1-65535)	10 sec
Link Trace Cache	☒ Enabled
Link Trace Cache Hold Time (1-65535)	100 min
Link Trace Cache Size (1-4095)	100 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

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 307: Configuring Interfaces for CFM



Administration > CFM	
Step: 2. Configure Interface	
Interface: ☒ Port ☐ Trunk	
Port List Total: 28	
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 32: 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 33: 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 308: 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 309: Showing Maintenance Domains

Administration > CFM

Step: 1. Configure MD Action: Show

CFM MD List Total: 1

	MD Index	MD Name	MD Level	MIP Creation Type
☐	1	voip	3	Explicit

Delete Revert

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 310: Configuring Detailed Settings for Maintenance Domains

The screenshot shows a web interface for configuring CFM (Connectivity Fault Management) settings. The breadcrumb navigation at the top reads "Administration > CFM". Below this, there are two dropdown menus: "Step: 3. Configure MD" and "Action: Configure Details". The main configuration area contains five rows of settings, each with a label and a corresponding input field:

Label	Value
MD Index	1
MEP Archive Hold Time (1-65535)	100
MEP Fault Notify Lowest Priority (1-6)	2
MEP Fault Notify Alarm Time (3-10)	3
MEP Fault Notify Reset Time (3-10)	10

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

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 496](#)).
- ◆ An MA must be defined before any associated DSAPs or remote MEPs can be assigned (see [“Configuring Remote Maintenance End Points” on page 498](#)).
- ◆ Multiple domains at the same maintenance level cannot have an MA on the same VLAN (see [“Configuring CFM Maintenance Domains” on page 487](#)).
- ◆ Before removing an MA, first remove the MEPs assigned to it (see [“Configuring Maintenance End Points” on page 496](#)).
- ◆ For a detailed description of the MIP types, refer to the Command Usage section under [“Configuring CFM Maintenance Domains” on page 487](#).

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)

13. The total length of the MD name and MA name cannot exceed 44 characters.

- ◆ **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.

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 311: Creating Maintenance Associations

Administration > CFM

Step: 4. Configure MA Action: Add

MD Index 1

MA Index (1-2147483647) 1

MA Name rd

Primary VLAN (1-4094) 1

MIP Creation Type Default

Apply 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 312: Showing Maintenance Associations

Administration > CFM

Step: 4. Configure MA Action: Show

MD Index 1

CFM MA List Total: 1

	MA Index	MA Name	Primary VLAN	MIP Creation Type
☐	1	rd	1	Default

Delete 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 313: Configuring Detailed Settings for Maintenance Associations

The screenshot shows the 'Administration > CFM' configuration page. At the top, there is a breadcrumb 'Administration > CFM'. Below it, a 'Step:' dropdown is set to '4. Configure MA' and an 'Action:' dropdown is set to 'Configure Details'. The main configuration area contains the following fields and controls:

MD Index	1
MA Index	1
MA Name Format	Character String
Interval Level (4-7)	4
Connectivity Check	☒ Enabled
Cross Check	☒ Enabled
AIS Status	☒ Enabled
AIS Period	1
AIS Transmit Level (0-7)	0
AIS Suppress Alarm	☐ Enabled

At the bottom right of the form 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 314: Configuring Maintenance End Points

The screenshot shows a web interface titled "Administration > CFM". Below the title bar, there are two dropdown menus: "Step:" set to "3. Configure MEP" and "Action:" set to "Add". The main form area contains the following fields:

- MD Index:** A dropdown menu with the value "1" selected.
- MA Index:** A dropdown menu with the value "1" selected.
- MEP ID (1-8191):** A text input field containing the value "1".
- MEP Direction:** A checkbox labeled "Up" which is currently unchecked.
- Interface:** A group of controls including a radio button selected for "Port" with a dropdown menu showing "1", and a radio button for "Trunk" with a dropdown menu.

At the bottom right of the form, there are two buttons: "Apply" and "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 315: Showing Maintenance End Points

Administration > CFM

Step: 3. Configure MEP Action: Show

MD Index: 1

MA Index: 1

CFM Maintenance Association End Point List Total: 1

	MEP ID	MEP Direction	Interface
☐	1	Down	Unit 1 / Port 1

Delete Revert

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 316: Configuring Remote Maintenance End Points

Administration > CFM

Step: 4. Configure Remote MEP Action: Add

MD Index 1

MA Index 1

MEP ID (1-8191) 2

Apply Revert

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 317: Showing Remote Maintenance End Points

Administration > CFM

Step: 4. Configure Remote MEP Action: Show

MD Index: 1
MA Index: 1

CFM Remote Maintenance Association End Point List Total: 1

	MEP ID
☐	2

Delete Revert

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 link trace 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 318: Transmitting Link Trace Messages

The screenshot shows a web interface for 'Administration > CFM'. The 'Step' dropdown is set to '7. Transmit Link Trace'. The form contains the following fields and values:

- MD Index: 1
- MA Index: 1
- Source MEP ID (1-8191): 1
- Target: ☒ MEP ID (1-8191) 2
- ☐ MAC Address
- TTL (0-255): 5

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

Transmitting Loopback 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 319: Transmitting Loopback Messages



Administration > CFM

Step: 8. Transmit Loopback

MD Index: 1

MA Index: 1

Source MEP ID (1-8191):

Target: ☒ MEP ID (1-8191)

☐ MAC Address

Counts (1-1024):

Packet Size (64-1518): bytes

Test

Result

```

Sending 200 Ethernet CFM loopback message, timeout is 5 sec.
Received 0 Ethernet CFM loopback message in 1 sec.
Received 0 Ethernet CFM loopback message in 5 secs. Success rate is 0% (0/5).
  
```

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 TxTimestampb (Timestamp at the time of transmitting a frame with DM reply information):
$$\text{Frame Delay} = (\text{RxTimeStampb} - \text{TxTimeStampf}) - (\text{TxTimestampb} - \text{RxTimeStampf})$$
- ◆ 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 320: 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 321: Showing Information on Local MEPs

Administration > CFM

Step: 7. 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	Eth 1/1	Enabled	00-E0-0C-00-00-FE

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 Character String or ICC Based.
- ◆ **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 322: 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 323: Showing Information on Local MIPs

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
voip	3	rd	1	Unit 1 / Port 7
voip	3	rd	1	Unit 1 / Port 8
voip	3	rd	1	Unit 1 / Port 9
voip	3	rd	1	Unit 1 / Port 10

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 324: Showing Information on Remote MEPs

CFM Remote Maintenance Association End Point Information Total: 1					
MEP ID	MA Name	Level	Primary VLAN	MEP Up	Remote MAC Address
1	ma1	0	2	Yes	00-01-02-03-04-05

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 325: 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
<input type="button" value="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 Address / 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 326: Showing the Link Trace Cache

Administration > CFM

Step: 8. Show Information Action: Show Link Trace Cache

CFM Link Trace Cache Information Total: 2

Hops	MA	IP/Alias Address	Forwarded	Ingress MAC Address	Egress MAC Address	Ingress Action	Egress Action	Reply
1	denver	11-11-11-11-11-11	Forwarded	11-11-11-11-11-12	11-11-11-11-11-13	OK	OK	FDB
2	boston	AA-BB-CC-00-04-00	Not Forwarded					

Clear

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.

¹⁴. See “Configuring CFM Maintenance Domains” on page 487.

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 327: Showing Settings for the 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	aa	bb	NONE	allDef	3	3

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.
- ◆ **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.

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.

- **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 328: Showing Continuity Check Errors

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

OAM Configuration

The switch provides OAM (Operation, Administration, and Maintenance) remote management tools required to monitor and maintain the links to subscriber CPEs (Customer Premise Equipment). This section describes functions including enabling OAM for selected ports, loopback testing, and displaying remote device information.

Enabling OAM on Local Ports

Use the Administration > OAM > Interface page to enable OAM functionality on the selected port. Not all CPEs support operation and maintenance functions, so OAM is therefore disabled by default. If a CPE supports OAM, this functionality must first be enabled on the connected port to gain access to the configuration functions provided under the OAM menu.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28)

- ◆ **Admin Status** – Enables or disables OAM functions. (Default: Disabled)
- ◆ **Operation State** – Shows the operational state between the local and remote OAM devices. This value is always “disabled” if OAM is disabled on the local interface.

Table 34: OAM Operation State

State	Description
Disabled	OAM is disabled on this interface via the OAM Admin Status.
Link Fault	The link has detected a fault or the interface is not operational.
Passive Wait	This value is returned only by OAM entities in passive mode and indicates the OAM entity is waiting to see if the peer device is OAM capable.
Active Send Local	This value is used by active mode devices and indicates the OAM entity is actively trying to discover whether the peer has OAM capability but has not yet made that determination.
Send Local And Remote	The local OAM entity has discovered the peer but has not yet accepted or rejected the configuration of the peer.
Send Local And Remote OK	OAM peering is allowed by the local device.
OAM Peering Locally Rejected	The local OAM entity rejects the peering.
OAM Peering Remotely Rejected	The remote OAM entity rejects the peering.
Operational	When the local OAM entity learns that both it and the remote OAM entity have accepted the peering, the state moves to this state.
Non Oper Half Duplex	This state is returned whenever Ethernet OAM is enabled but the interface is in half-duplex operation.

- ◆ **Mode** – Sets the OAM operation mode. (Default: Active)
 - **Active** – All OAM functions are enabled.
 - **Passive** – All OAM functions are enabled, except for OAM discovery, sending variable request OAMPDUs, and sending loopback control OAMPDUs.
- ◆ **Critical Link Event** – Controls reporting of critical link events to its OAM peer.
 - **Dying Gasp** – If an unrecoverable condition occurs, the local OAM entity (i.e., this switch) indicates this by immediately sending a trap message. (Default: Enabled)

Dying gasp events are caused by an unrecoverable failure, such as a power failure or device reset.



Note: When system power fails, the switch will always send a dying gasp trap message prior to power down.

- **Critical Event** – If a critical event occurs, the local OAM entity indicates this to its peer by setting the appropriate flag in the next OAMPDU to be sent and stores this information in its OAM event log. (Default: Enabled)

Critical events include various failures, such as abnormal voltage fluctuations, out-of-range temperature detected, fan failure, CRC error in flash memory, insufficient memory, or other hardware faults.

- ◆ **Errored Frame** – Controls reporting of errored frame link events.

An errored frame is a frame in which one or more bits are errored.

An errored frame link event occurs if the threshold is reached or exceeded within the specified period.

If reporting is enabled and an errored frame link event occurs, the local OAM entity (this switch) sends an Event Notification OAMPDU to the remote OAM entity. The Errored Frame Event TLV includes the number of errored frames detected during the specified period.

- **Status** – Enables reporting of errored frame link events. (Default: Enabled)
- **Window Size** – The period of time in which to check the reporting threshold for errored frame link events. (Range: 10-65535 in units of 10 milliseconds; Default: 10 units of 10 milliseconds, or the equivalent of 1 second)
- **Threshold Count** – The threshold for errored frame link events. (Range: 1-65535; Default: 1)

Web Interface

To enable OAM functionality on the selected port:

1. Click Administration, OAM, Interface.
2. Set the OAM administrative status and operational mode for the required ports. Specify whether or not critical link events will be reported by the switch. Specify whether errored frame link events will be reported, as well as the required window size and threshold.
3. Click Apply.

Figure 329: Enabling OAM for Local Ports

Administration > OAM > Interface

OAM Port List Total: 52

Port	Admin Status	Operation State	Mode	Critical Link Event		Errored Frame		
				Dying Gasp	Critical Event	Status	Window Size (10-65535 1/10 sec)	Threshold Count (1-65535)
1	☐ Enabled	Link Fault	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
2	☐ Enabled	Link Fault	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
3	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
4	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
5	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
6	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
7	☐ Enabled	Link Fault	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
8	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
9	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
10	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1

Apply Revert

Displaying Statistics for OAM Messages Use the Administration > OAM > Counters page to display statistics for the various types of OAM messages passed across each port.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28)
- ◆ **Clear** – Clears statistical counters for the selected ports.
- ◆ **OAMPDU** – Message types transmitted and received by the OAM protocol, including Information OAMPDUs, unique Event OAMPDUs, Loopback Control OAMPDUs, and Organization Specific OAMPDUs.

Web Interface

To display statistics for OAM messages:

1. Click Administration, OAM, Counters.

Figure 330: Displaying Statistics for OAM Messages

Administration > OAM > Counters

OAM Port Counters Total: 52

123456

Port	OAMPDU							
	Information		Event Notification		Loopback Control		Organization Specific	
	Transmitted	Received	Transmitted	Received	Transmitted	Received	Transmitted	Received
1	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0

Clear

Displaying the OAM Event Log Use the Administration > OAM > Event Log page to display link events for the selected port.

Command Usage

- ◆ When a link event occurs, no matter whether the location is local or remote, this information is entered in OAM event log.
- ◆ When the log system becomes full, older events are automatically deleted to make room for new entries.
- ◆ The time of locally generated events can be accurately retrieved from the sysUpTime variable. For remotely generated events, the time of an event is indicated by the reception of an Event Notification OAMPDU from the peer.

Web Interface

To display link events for the selected port:

1. Click Administration, OAM, Event Log.
2. Select a port from the drop-down list.

Figure 331: Displaying the OAM Event Log



Displaying the Status of Remote Interfaces Use the Administration > OAM > Remote Interface page to display information about attached OAM-enabled devices.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28)
- ◆ **MAC Address** – MAC address of the OAM peer.
- ◆ **OUI** – Organizational Unit Identifier of the OAM peer.
- ◆ **Remote Loopback** – Shows if remote loopback is supported by the OAM peer.
- ◆ **Unidirectional Function** – Shows if this function is supported by the OAM peer.

If supported, this indicates that the OAM entity supports the transmission of OAMPDUs on links that are operating in unidirectional mode (where traffic flows in one direction only). Some newer physical layer devices support the optional ability to encode and transmit data while one direction of the link is non-operational. This function allows OAM remote fault indication during fault conditions. This switch does not support the unidirectional function, but can parse error messages sent from a peer with unidirectional capability.
- ◆ **Link Monitor** – Shows if the OAM entity can send and receive Event Notification OAMPDUs.
- ◆ **MIB Variable Retrieval** – Shows if the OAM entity can send and receive Variable Request and Response OAMPDUs.

Web Interface

To display information about attached OAM-enabled devices:

1. Click Administration, OAM, Remote Interface.

Figure 332: Displaying Status of Remote Interfaces

Port	MAC Address	OUI	Remote Loopback	Unidirectional Function	Link Monitor	MIB Variable Retrieval
1	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
2	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
3	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
4	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
5	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
6	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
7	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
8	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
9	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
10	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled

Configuring a Remote Loopback Test

Use the Administration > OAM > Remote Loopback (Remote Loopback Test) page to initiate a loop back test to the peer device attached to the selected port.

Command Usage

- ◆ You can use this command to perform an OAM remote loop back test on the specified port. The port that you specify to run this test must be connected to a peer OAM device capable of entering into OAM remote loop back mode.
- ◆ During a remote loop back test, the remote OAM entity loops back every frame except for OAMPDUs and pause frames.
- ◆ OAM remote loopback can be used for fault localization and link performance testing. Statistics from both the local and remote DTE can be queried and compared at any time during loop back testing.
- ◆ To perform a loopback test, first enable Remote Loop Back Mode, click Test, and then click End. The number of packets transmitted and received will be displayed.

Parameters

These parameters are displayed:

Loopback Mode of Remote Device

- ◆ **Port** – Port identifier. (Range: 1-28)
- ◆ **Loopback Mode** – Shows if loop back mode is enabled on the peer. This attribute must be enabled before starting the loopback test.

- ◆ **Loopback Status** – Shows if loopback testing is currently running.

Loopback Test Parameters

- ◆ **Packet Number** – Number of packets to send. (Range: 1-99999999; Default: 10000)
- ◆ **Packet Size** – Size of packets to send. (Range: 64-1518 bytes; Default: 64 bytes)
- ◆ **Test** – Starts the loop back test.
- ◆ **End** – Stops the loop back test.

Loop Back Status of Remote Device

- ◆ **Result** – Shows the loop back status on the peer. The loop back states shown in this field are described below.

Table 35: Remote Loopback Status

State	Description
No Loopback	Operating in normal mode with no loopback in progress.
Initiating Loopback	The local OAM entity is starting the loopback process with its peer. It has yet to receive any acknowledgement that the remote OAM entity has received its loopback command request.
Remote Loopback	The local OAM client knows that the remote OAM entity is in loopback mode.
Terminating Loopback	The local OAM client is in the process of terminating the remote loopback.
Local Loopback	The remote OAM client has put the local OAM entity in loopback mode.
Unknown	This status may be returned if the OAM loopback is in a transition state but should not persist.

- **Packets Transmitted** – The number of loop back frames transmitted during the last loopback test on this interface.
- **Packets Received** – The number of loop back frames received during the last loopback test on this interface.
- **Loss Rate** – The percentage of packets for which there was no response.

Web Interface

To initiate a loop back test to the peer device attached to the selected port:

1. Click Administration, OAM, Remote Loop Back.
2. Select Remote Loopback Test from the Action list.

3. Select the port on which to initiate remote loop back testing, enable the Loop Back Mode attribute, and click Apply.
4. Set the number of packets to send and the packet size, and then click Test.

Figure 333: Running a Remote Loop Back Test

Administration > OAM > Remote Loopback

Action: Remote Loopback Test

Port: 2

Loopback Mode: ☒ Enabled

Loopback Status: Remote Loopback

Apply Revert

Packet Number (1-99999999):

Packet Size (64-1518): bytes

Test End

Result

10 packets transmitted. 10 packets received. Loss rate is 0.00 %

Displaying Results of Remote Loopback Testing

Use the Administration > OAM > Remote Loopback (Show Test Result) page to display the results of remote loop back testing for each port for which this information is available.

Parameters


These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28)
- ◆ **Packets Transmitted** – The number of loop back frames transmitted during the last loop back test on this interface.
- ◆ **Packets Received** – The number of loop back frames received during the last loop back test on this interface.
- ◆ **Loss Rate** – The percentage of packets transmitted for which there was no response.

Web Interface

To display the results of remote loop back testing for each port for which this information is available:

1. Click Administration, OAM, Remote Loopback.
2. Select Show Test Result from the Action list.

Figure 334: Displaying the Results of Remote Loop Back Testing

Administration > OAM > Remote Loopback

Action:

Port Remote Test Result List Total: 28

Port	Packets Transmitted	Packets Received	Loss Rate
1	0	0	0.00 %
2	0	0	0.00 %
3	0	0	0.00 %
4	0	0	0.00 %
5	0	0	0.00 %
6	0	0	0.00 %
7	0	0	0.00 %
8	0	0	0.00 %
9	0	0	0.00 %
10	0	0	0.00 %

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. M_{slow} 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 335: Configuring UDLD Protocol Intervals

Administration > UDLD

Step: 1. Configure Global

Message Interval (7-90)	15	seconds
Detection Interval (5-255)	5	seconds
Recovery Status	☐ Enabled	
Recovery Interval (30-86400)	300	seconds

Apply 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)
- ◆ **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 336: Configuring UDLD Interface Settings



Port	UDLD	Aggressive Mode	Operation State	Port State	Message Interval (seconds)	Detection Interval (seconds)
1	☒ Enabled	☒ Enabled	Disabled	Unknown	7	5
2	☒ Enabled	☒ Enabled	Disabled	Unknown	7	5
3	☒ Enabled	☒ Enabled	Disabled	Unknown	7	5
4	☒ Enabled	☒ Enabled	Disabled	Unknown	7	5
5	☒ Enabled	☒ 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: 1-28)
- ◆ **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 337: Displaying UDLD Neighbor Information



Administration > UDLD

Step: 3. Show Information

Port 1

UDLD Neighbor List Total: 3

Entry	Device ID	Port ID	Device Name	Neighbor State	Expire (seconds)	Message Interval (seconds)	Detection Interval (seconds)
1	SQ-0	Eth 1/1	GTL-2891	Bidirectional	40	15	5
2	SQ-1	Eth 1/1	GTL-2891	Unidirectional	39	14	4
3	SQ-2	Eth 1/1	GTL-2891	Unidirectional	38	13	3

Multicast Filtering


This chapter describes how to configure the following multicast services:

- ◆ [IGMP Snooping](#) – Configures snooping and query parameters.
- ◆ [Filtering and Throttling](#) – Filters specified multicast service, or throttles the maximum of multicast groups allowed on an interface.
- ◆ [MLD Snooping](#) – Configures snooping and query parameters for IPv6.
- ◆ [Layer 3 IGMP](#) – Configures IGMP query used with multicast routing.
- ◆ [Multicast VLAN Registration for IPv4](#) – Configures a single network-wide multicast VLAN shared by hosts residing in other standard or private VLAN groups, preserving security and data isolation.
- ◆ [Multicast VLAN Registration for IPv6](#) – Configures a single network-wide multicast VLAN shared by hosts residing in other standard or private VLAN groups, preserving security and data isolation.

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 338: 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 753](#) and [769](#)).

You can also configure a single network-wide multicast VLAN shared by hosts residing in other standard or private VLAN groups, preserving security and data isolation [“Multicast VLAN Registration for IPv4” on page 579](#).

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 339: 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 535](#)) 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 a Multicast Router” on page 539](#)). 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 1023 multicast entries can be maintained for IGMP snooping and 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 535](#)).

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 539](#)). 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 541](#)).

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 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 DVMRP or 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: Enabled)

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 544](#)).

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, Recommended Range: 300-500 seconds, Default: 300)

- ◆ **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 340: Configuring General Settings for IGMP Snooping

Multicast > IGMP Snooping > General

IGMP Snooping Status	☒ Enabled
Proxy Reporting Status	☐ Enabled
TCN Flood	☐ Enabled
TCN Query Solicit	☐ Enabled
Router Alert Option	☐ Enabled
Unregistered Data Flooding	☐ Enabled
Forwarding Priority (0-7)	☐ <input type="text"/>
Version Exclusive	☐ 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	☐ Enabled

Apply Revert

Specifying Static Interfaces for a Multicast Router

Use the Multicast > IGMP Snooping > Multicast Router (Add Static Multicast Router) page to statically attach an 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 535](#)) 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 341: Configuring a Static Interface for a Multicast Router

The screenshot shows a web interface titled "Multicast > IGMP Snooping > Multicast Router". Below the title bar, there is a section labeled "Action:" with a dropdown menu set to "Add Static Multicast Router". Below this, there is a "VLAN" field with a dropdown menu set to "1". Underneath the VLAN field, there is an "Interface" section with two radio buttons: "Port" (which is selected) and "Trunk". The "Port" radio button is accompanied by a dropdown menu set to "1". The "Trunk" radio button is accompanied by a dropdown menu. At the bottom right of the form, there are two buttons: "Apply" and "Revert".

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 342: Showing Static Interfaces Attached a Multicast Router

Multicast > IGMP Snooping > Multicast Router

Action: Show Static Multicast Router

VLAN: 1

Static Multicast Router Interface List Total: 6

	Interface
☐	Unit 1 / Port 1
☐	Unit 1 / Port 2
☐	Unit 1 / Port 3
☐	Trunk 2
☐	Trunk 5
☐	Unit 1 / Port 4

Delete Revert

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 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 343: Showing Current Interfaces Attached a Multicast Router

Multicast > IGMP Snooping > Multicast Router

Action: Show Current Multicast Router

VLAN: 1

Multicast Router Interface Information Total: 2

Interface	Type	Expire
Eth 1 / 1	Static	
Eth 1 / 4	Dynamic	0:4:59

Assigning Interfaces to Multicast Services

Use the Multicast > IGMP Snooping > IGMP Member (Add Static Member) page to statically assign a 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 535](#)). 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 a 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 344: Assigning an Interface to a Multicast Service

The screenshot shows a web interface for configuring IGMP Member settings. The breadcrumb trail at the top reads "Multicast > IGMP Snooping > IGMP Member". Below this, there is a section for "Action:" with a dropdown menu set to "Add Static Member". The main configuration area contains three rows: "VLAN" with a dropdown menu set to "1", "Interface" with radio buttons for "Port" (selected) and "Trunk", and "Multicast IP" with a text input field containing "224.1.1.1". At the bottom right, there are two buttons: "Apply" and "Revert".

To show the static interfaces assigned to a 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 345: Showing Static Interfaces Assigned to a Multicast Service




The screenshot shows the 'Multicast > IGMP Snooping > IGMP Member' configuration page. The 'Action' dropdown is set to 'Show Static Member'. The 'VLAN' dropdown is set to '1'. Below this, the 'IGMP Member Interface List' is displayed with a 'Total: 6' count. The table lists six interfaces with their corresponding multicast IP addresses. At the bottom, there are 'Delete' and 'Revert' buttons.

	Interface	Multicast IP
☐	Unit 1 / Port 1	224.1.1.1
☐	Unit 1 / Port 2	224.1.2.2
☐	Unit 1 / Port 3	230.1.1.1
☐	Trunk 2	230.1.2.2
☐	Trunk 5	239.1.1.1
☐	Unit 1 / Port 4	239.2.2.2

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 346: Showing Current Interfaces Attached a Multicast Router



The screenshot shows the 'Multicast > IGMP Snooping > Multicast Router' configuration page. The 'Action' dropdown is set to 'Show Current Multicast Router'. The 'VLAN' dropdown is set to '1'. Below this, the 'Multicast Router Interface Information' is displayed with a 'Total: 2' count. The table lists two interfaces with their types and expiration times. At the bottom, there are 'Delete' and 'Revert' buttons.

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 535](#).

Command Usage

Multicast Router Discovery

There have been many mechanisms used in the past to identify multicast routers. This has led 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 an 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: Enabled)

When IGMP snooping is enabled globally (see [page 535](#)), 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 or IGMPv3 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)

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.

- ◆ **Query Interval** – The interval between sending IGMP proxy general queries. (Range: 2-31744 seconds; Default: 125 seconds)

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 535](#)), or as a proxy host when IGMP snooping proxy reporting is enabled ([page 535](#)).

- ◆ **Query Response Interval** – The maximum time the system waits for a response to 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 535](#)), or as a proxy host when IGMP snooping proxy reporting is enabled ([page 535](#)).

- ◆ **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 ([page 535](#)) or IGMP querier is enabled ([page 535](#)).

- ◆ **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 347: 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): 125 seconds

Query Response Interval (10-31740): 100 (1/10 seconds, multiple of 10)

Last Member Query Interval (1-31744): 10 (1/10 seconds, multiple of 10)

Last Member Query Count (1-255): 2

Proxy (Query) Address: 0.0.0.0

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 348: Showing Interface Settings for IGMP Snooping

Multicast > IGMP Snooping > Interface

Action: Show VLAN Information

IGMP Snooping VLAN List Total: 4

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 and Multicast Data

Use the Multicast > IGMP Snooping > Interface (Configure Interface) page to configure an interface to drop IGMP query packets or multicast data 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.
- ◆ **Multicast Data Drop** – Configures an interface to stop multicast services from being forwarded to users attached to the downstream port (i.e., the interfaces specified by this command).

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. Select Port or Trunk interface.
4. Enable the required drop functions for any interface.
5. Click Apply.

Figure 349: Dropping IGMP Query or Multicast Data Packets

The screenshot shows the 'Multicast > IGMP Snooping > Interface' configuration page. At the top, there is a breadcrumb trail and a title bar. Below the title bar, there is an 'Action:' dropdown menu set to 'Configure Interface'. Underneath, there are radio buttons for 'Interface', 'Port' (selected), and 'Trunk'. A 'Port List' section shows a table with 5 ports. Each port has two checkboxes: 'IGMP Query Drop' and 'Multicast Data Drop'. Both checkboxes are currently unchecked, indicating they are disabled. The table has a 'Total: 28' label and three numbered tabs (1, 2, 3) on the right side.

Port	IGMP Query Drop	Multicast Data Drop
1	☐ Enabled	☐ Enabled
2	☐ Enabled	☐ Enabled
3	☐ Enabled	☐ Enabled
4	☐ Enabled	☐ Enabled
5	☐ Enabled	☐ Enabled

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 535](#)).

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.
- ◆ **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** – 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 350: Showing Multicast Groups Learned by IGMP Snooping

Multicast > IGMP Snooping > Forwarding Entry

VLAN:

IGMP Snooping Forwarding Entry List Total: 10

Group 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	Eth 1 / 9 (Router Port)		02:24	1 (Port)
224.1.1.3	Eth 1 / 9 (Router Port)	00:00:16:14		1 (Port)
	Eth 1 / 9 (Router 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)

Clear Click this button to clear all IGMP Snooping dynamic groups.

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: 1-28)
- ◆ **Trunk** – Trunk identifier. (Range: 1-16)

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 351: Displaying IGMP Snooping Statistics – Query

Multicast > IGMP Snooping > Statistics

Action: Show Query Statistics

VLAN 1

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

Clear All [Click this button to clear all IGMP Snooping statistics.](#)

Refresh

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 352: Displaying IGMP Snooping Statistics – VLAN

Multicast > IGMP Snooping > Statistics

Action: Show VLAN Statistics ▾

VLAN 1 ▾

Input Statistics

Report	0	Drop	0
Leave	0	Join Success	0
G Query	0	Group	0
G(-S)-S Query	0		

Output Statistics

Report	18
Leave	0
G Query	0
G(-S)-S Query	0

Refresh

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 353: Displaying IGMP Snooping Statistics – Port

Multicast > IGMP Snooping > Statistics

Action: Show Port Statistics ▾

Port 1 ▾

Input Statistics

Report	0	Drop	0
Leave	0	Join Success	0
G Query	0	Group	0
G(-S)-S Query	0		

Output Statistics

Report	0
Leave	0
G Query	0
G(-S)-S Query	0

Refresh

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 354: Enabling IGMP Filtering and Throttling

Multicast > IGMP Snooping > Filter

Step: 1. Configure General

IGMP Filter Status ☒ Enabled

Apply Revert

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 355: Creating an IGMP Filtering Profile

Multicast > IGMP Snooping > Filter

Step: 2. Configure Profile Action: Add

Profile ID (1-4294967295) 19

Access Mode Permit

Apply Revert

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 356: Showing the IGMP Filtering Profiles Created

Multicast > IGMP Snooping > Filter

Step: 2. Configure Profile Action: Show

IGMP Snooping Filter Profile List Total: 1

	Profile ID	Action Mode
☐	19	Permit

Delete Revert

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 357: Adding Multicast Groups to an IGMP Filtering Profile

Multicast > IGMP Snooping > Filter

Step: 2. Configure Profile Action: Add Multicast Group Range

Profile ID: 19

Start Multicast IP Address: 239.2.3.1

End Multicast IP Address: 239.2.3.200

Apply Revert

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 358: Showing the Groups Assigned to an IGMP Filtering Profile

Multicast > IGMP Snooping > Filter

Step: 2. Configure Profile Action: Show Multicast Group Range

Profile ID: 19

Multicast IP Address Range List Total: 1

	Start Multicast IP Address	End Multicast IP Address
☐	239.2.3.1	239.2.3.200

Delete Revert

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 359: Configuring IGMP Filtering and Throttling Interface Settings

Multicast > IGMP Snooping > Filter

Step: 3. Configure Interface

Interface ☒ Port ☐ Trunk

IGMP Filter and Throttling Port List Total: 28

Port	Profile ID	Max Multicast Groups (1-1024)	Current Multicast Groups	Throttling Action Mode	Throttling Status
1	19	64	0	Deny	False
2	(none)	1024	0	Deny	False
3	(none)	1024	0	Deny	False
4	(none)	1024	0	Deny	False
5	(none)	1024	0	Deny	False

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 360: Configuring General Settings for MLD Snooping

Multicast > MLD Snooping > General

MLD Snooping Status	☐ Enabled
Querier Status	☐ Enabled
Robustness (2-10)	<input type="text" value="2"/>
Query Interval (60-125)	<input type="text" value="125"/> seconds
Query Max Response Time (5-25)	<input type="text" value="10"/> seconds
Router Port Expiry Time (300-500)	<input type="text" value="300"/> seconds
MLD Snooping Version (1-2)	<input type="text" value="2"/>
Unknown Multicast Mode	<input type="text" value="To Router Port"/>

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 361: Configuring Immediate Leave for MLD Snooping

Multicast > MLD Snooping > Interface

VLAN 1

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 561](#)) 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 362: Configuring a Static Interface for an IPv6 Multicast Router

Multicast > MLD Snooping > Multicast Router

Action: Add Static Multicast Router

VLAN: 1

Interface: ☒ Port 1 ☐ Trunk

Apply Revert

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 363: Showing Static Interfaces Attached an IPv6 Multicast Router

Multicast > MLD Snooping > Multicast Router

Action: Show Static Multicast Router

VLAN: 1

Static Multicast Router Interface List Total: 2

Interface
Unit 1 / Port 1
Unit 1 / Port 2

Delete Revert

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 364: Showing Current Interfaces Attached an IPv6 Multicast Router

Multicast > MLD Snooping > Multicast Router

Action: Show Current Multicast Router

VLAN: 1

Multicast Router Interface Information Total: 4

Interface	Type
Unit 1 / Port 4	Static
Unit 1 / Port 5	Dynamic
Trunk 2	Dynamic
Trunk 3	Dynamic

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 561](#)). 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 365: 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 366: 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
☐	FF02::01:01:01:01	Unit 1 / Port 1
☐	FF02::01:01:01:02	Unit 1 / Port 2
☐	FF01::1	Unit 1 / Port 12
☐	FF01::2	Unit 1 / Port 13
☐	FF01::3	Unit 1 / Port 14
☐	FF01::4	Unit 1 / Port 15
☐	FF01::5	Unit 1 / Port 16
☐	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 367: Showing Current Interfaces Assigned to an IPv6 Multicast Service

Multicast > MLD Snooping > MLD Member

Action: Show Current Member

VLAN: 1

MLD Snooping Member Interface List Total: 9

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 368: Showing IPv6 Multicast Services and Corresponding Sources



Layer 3 IGMP (Query used with Multicast Routing)

IGMP Snooping – IGMP Snooping ([page 535](#)) 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 573](#)).



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 544](#).)


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 369: 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 748](#)).
2. Enable IGMP on the downstream interfaces which require proxy multicast service (see [“Configuring IGMP Interface Parameters” on page 573](#)).
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 370: Configuring IGMP Proxy Routing

Multicast > IGMP > Proxy

VLAN	1
IGMP Proxy Status	☒ Enabled
Unsolicited Report Interval (1-65535)	400 seconds

Apply Revert

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 748](#)).

Parameters

These parameters are displayed:

- ◆ **VLAN** – VLAN interface bound to a primary IP address. (Range: 1-4094)
- ◆ **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 371: Configuring IGMP Interface Settings

Multicast > IGMP > Interface

VLAN	1	
IGMP Protocol Status	☒ 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 372: 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 373: Showing Static IGMP Groups

Multicast > IGMP > Static Group

Action: Show

VLAN: 1

IGMP Static Group List Total: 1

	Static Group Address	Source Address
☐	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 551](#)).

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 573](#)), and multicast routing must be enabled globally on the system (see [“Configuring Global Settings for Multicast Routing” on page 748](#)).

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 374: 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 375: Displaying Multicast Groups Learned from IGMP (Details)

Multicast > IGMP > Group Information				
Action: Show Detail				
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	10:20:00	YES	
11.2.2.2	0h:1m:7s	10:20:00	YES	
12.2.2.2	0h:1m:7s	10:20:00	NO	


Multicast VLAN Registration for IPv4

Multicast VLAN Registration (MVR) is a protocol that controls access to a single network-wide VLAN most commonly used for transmitting multicast traffic (such as television channels or video-on-demand) across a service provider's network. Any multicast traffic entering an MVR VLAN is sent to all attached subscribers. This protocol can significantly reduce to processing overhead required to dynamically monitor and establish the distribution tree for a normal multicast VLAN. This makes it possible to support common multicast services over a wide part of the network without having to use any multicast routing protocol.

MVR maintains the user isolation and data security provided by VLAN segregation by passing only multicast traffic into other VLANs to which the subscribers belong. Even though common multicast streams are passed onto different VLAN groups

from the MVR VLAN, users in different IEEE 802.1Q or private VLANs cannot exchange any information (except through upper-level routing services).

Figure 376: MVR Concept



Command Usage

◆ General Configuration Guidelines for MVR:

1. Enable MVR for a domain on the switch, and select the MVR VLAN (see [“Configuring MVR Domain Settings” on page 583](#)).
2. Create an MVR profile by specifying the multicast groups that will stream traffic to attached hosts, and assign the profile to an MVR domain (see [“Configuring MVR Group Address Profiles” on page 584](#)).
3. Set the interfaces that will join the MVR as source ports or receiver ports (see [“Configuring MVR Interface Status” on page 587](#)).
4. For multicast streams that will run for a long term and be associated with a stable set of hosts, you can statically bind the multicast group to the participating interfaces (see [“Assigning Static MVR Multicast Groups to Interfaces” on page 589](#)).

- ◆ Although MVR operates on the underlying mechanism of IGMP snooping, the two features operate independently of each other. One can be enabled or disabled without affecting the behavior of the other. However, if IGMP snooping and MVR are both enabled, MVR reacts only to join and leave messages from multicast groups configured under MVR. Join and leave messages from all other multicast groups are managed by IGMP snooping. Also, note that only IGMP version 2 or 3 hosts can issue multicast join or leave messages. Since IGMP version 1 hosts do not support leave messages, they are timed out by the switch.

Configuring MVR Global Settings Use the Multicast > MVR (Configure Global) page to configure proxy switching and the robustness variable.

Parameters

These parameters are displayed:

- ◆ **Proxy Switching** – Configures MVR proxy switching, where the source port acts as a host, and the receiver port acts as an MVR router with querier service enabled. (Default: Enabled)
 - When MVR proxy-switching is enabled, an MVR source port serves as the upstream or host interface, and the MVR receiver port serves as the querier. The source port performs only the host portion of MVR by sending summarized membership reports, and automatically disables MVR router functions.
 - Receiver ports are known as downstream or router interfaces. These interfaces perform the standard MVR router functions by maintaining a database of all MVR subscriptions on the downstream interface. Receiver ports must therefore be configured on all downstream interfaces which require MVR proxy service.
 - When the source port receives report and leave messages, it only forwards them to other source ports.
 - When receiver ports receive any query messages, they are dropped.
 - When changes occurring in the downstream MVR groups are learned by the receiver ports through report and leave messages, an MVR state change report is created and sent to the upstream source port, which in turn forwards this information upstream.
 - When MVR proxy switching is disabled:
 - Any membership reports received from receiver/source ports are forwarded to all source ports.
 - When a source port receives a query message, it will be forwarded to all downstream receiver ports.
 - When a receiver port receives a query message, it will be dropped.
- ◆ **Robustness Value** – Configures the expected packet loss, and thereby the number of times to generate report and group-specific queries. (Range: 1-255; Default: 2)
 - This parameter is used to set the number of times report messages are sent upstream when changes are learned about downstream groups, and the number of times group-specific queries are sent to downstream receiver ports.
 - This parameter only takes effect when MVR proxy switching is enabled.

- ◆ **Proxy Query Interval** – Configures the interval at which the receiver port sends out general queries. (Range: 2-31744 seconds; Default: 125 seconds)
 - This parameter sets the general query interval at which active receiver ports send out general queries.
 - This interval is only effective when proxy switching is enabled.
- ◆ **Source Port Mode** – Configures the switch to forward any multicast streams within the parameters set by a profile, or to only forward multicast streams which the source port has dynamically joined.
 - **Always Forward** – By default, the switch forwards any multicast streams within the address range set by a profile, and bound to a domain. The multicast streams are sent to all source ports on the switch and to all receiver ports that have elected to receive data on that multicast address. (This is the default setting.)
 - **Dynamic** – When dynamic mode is enabled, the switch only forwards multicast streams which the source port has dynamically joined. In other words, both the receiver port and source port must subscribe to a multicast group before a multicast stream is forwarded to any attached client. Note that the requested streams are still restricted to the address range which has been specified in a profile and bound to a domain.

Web Interface

To configure global settings for MVR:

1. Click Multicast, MVR.
2. Select Configure Global from the Step list.
3. Set the status for MVR proxy switching, the robustness value used for report and query messages, the proxy query interval, and source port mode.
4. Click Apply.

Figure 377: Configuring Global Settings for MVR

The screenshot shows a web interface for configuring Multicast VLAN Registration (MVR) settings. The page title is "Multicast > MVR". Below the title, there is a "Step:" dropdown menu set to "1. Configure Global". The main configuration area contains four settings:

- Proxy Switching:** A checkbox labeled "Enabled" is checked.
- Robustness Value (1-255):** A text input field containing the value "2".
- Proxy Query Interval (2-31744):** A text input field containing the value "125", followed by the unit "sec".
- Source Port Mode:** A dropdown menu set to "Always Forward".

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

Configuring MVR Domain Settings Use the Multicast > MVR (Configure Domain) page to enable MVR globally on the switch, and select the VLAN that will serve as the sole channel for common multicast streams supported by the service provider.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **MVR Status** – When MVR is enabled on the switch, any multicast data associated with an MVR group is sent from all designated source ports, to all receiver ports that have registered to receive data from that multicast group. (Default: Disabled)
- ◆ **MVR VLAN** – Identifier of the VLAN that serves as the channel for streaming multicast services using MVR. MVR source ports should be configured as members of the MVR VLAN (see [“Adding Static Members to VLANs” on page 165](#)), but MVR receiver ports should not be manually configured as members of this VLAN. (Default: 1)
- ◆ **MVR Running Status** – Indicates whether or not all necessary conditions in the MVR environment are satisfied. Running status is Active as long as MVR is enabled, the specified MVR VLAN exists, and a source port with a valid link has been configured (see [“Configuring MVR Interface Status” on page 587](#)).
- ◆ **MVR Current Learned Groups** – The number of MVR groups currently assigned to this domain.
- ◆ **Forwarding Priority** – The CoS priority assigned to all multicast traffic forwarded into this domain. (Range: 0-7, where 7 is the highest priority)

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.
- ◆ **Upstream Source IP** – The source IP address assigned to all MVR control packets sent upstream on the specified domain. By default, all MVR reports sent upstream use a null source IP address.

Web Interface

To configure settings for an MVR domain:

1. Click Multicast, MVR.
2. Select Configure Domain from the Step list.
3. Select a domain from the scroll-down list.
4. Enable MVR for the selected domain, select the MVR VLAN, set the forwarding priority to be assigned to all ingress multicast traffic, and set the source IP address for all control packets sent upstream as required.

5. Click Apply.

Figure 378: Configuring Domain Settings for MVR

The screenshot shows the 'Multicast > MVR' configuration page. At the top, there is a 'Step:' dropdown menu set to '2. Configure Domain'. Below this, several configuration fields are listed: 'Domain ID' is a dropdown menu set to '1'; 'MVR Status' is a checkbox labeled 'Enabled' which is checked; 'MVR VLAN' is a dropdown menu set to '1'; 'MVR Running Status' is a text field containing 'Active'; 'MVR Current Learned Groups' is a text field containing '0'; 'Forwarding Priority (0-7)' is a checkbox and a text field, both currently empty; and 'Upstream Source IP' is a text field containing '192.168.0.12'. At the bottom right of the form are two buttons: 'Apply' and 'Revert'.

Configuring MVR Group Address Profiles

Use the Multicast > MVR (Configure Profile and Associate Profile) pages to assign the multicast group address for required services to one or more MVR domains.

Command Usage

- ◆ Use the Configure Profile page to statically configure all multicast group addresses that will join the MVR VLAN. Any multicast data associated with an MVR group is sent from all source ports to all receiver ports that have registered to receive data from that multicast group.
- ◆ The IP address range from 224.0.0.0 to 239.255.255.255 is used for multicast streams. MVR group addresses cannot fall within the reserved IP multicast address range of 224.0.0.x.
- ◆ IGMP snooping and MVR share a maximum number of 1024 groups. Any multicast streams received in excess of this limitation will be flooded to all ports in the associated domain.

Parameters

These parameters are displayed:

Configure Profile

- ◆ **Profile Name** – The name of a profile containing one or more MVR group addresses. (Range: 1-21 characters)
- ◆ **Start IP Address** – Starting IP address for an MVR multicast group. (Range: 224.0.1.0 - 239.255.255.255)
- ◆ **End IP Address** – Ending IP address for an MVR multicast group. (Range: 224.0.1.0 - 239.255.255.255)

Associate Profile

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Profile Name** – The name of a profile to be assigned to this domain. (Range: 1-21 characters)

Web Interface

To configure an MVR group address profile:

1. Click Multicast, MVR.
2. Select Configure Profile from the Step list.
3. Select Add from the Action list.
4. Enter the name of a group profile to be assigned to one or more domains, and specify a multicast group that will stream traffic to participating hosts.
5. Click Apply.

Figure 379: Configuring an MVR Group Address Profile

The screenshot shows the 'Multicast > MVR' configuration page. The 'Step' dropdown is set to '2. Configure Profile' and the 'Action' dropdown is set to 'Add'. The 'Profile Name' field contains 'sales', the 'Start IP Address' field contains '234.5.6.8', and the 'End IP Address' field contains '234.5.6.10'. There are 'Apply' and 'Revert' buttons at the bottom right.

To show the configured MVR group address profiles:

1. Click Multicast, MVR.
2. Select Configure Profile from the Step list.
3. Select Show from the Action list.

Figure 380: Displaying MVR Group Address Profiles

The screenshot shows the 'Multicast > MVR' configuration page. The 'Step' dropdown is set to '2. Configure Profile' and the 'Action' dropdown is set to 'Show'. Below the dropdowns, there is a table titled 'MVR Profile List' with a total of 2 profiles. The table has columns for 'Profile Name', 'Start IP Address', and 'End IP Address'. There are 'Delete' and 'Revert' buttons at the bottom right.

	Profile Name	Start IP Address	End IP Address
☐	sales	234.5.6.8	234.5.6.10
☐	rd	234.5.6.7	234.5.6.7

To assign an MVR group address profile to a domain:

1. Click Multicast, MVR.
2. Select Associate Profile from the Step list.
3. Select Add from the Action list.
4. Select a domain from the scroll-down list, and enter the name of a group profile.
5. Click Apply.

Figure 381: Assigning an MVR Group Address Profile to a Domain

Multicast > MVR

Step: 4. Associate Profile Action: Add

Domain ID: 1

Profile Name: sales

Apply Revert

To show the MVR group address profiles assigned to a domain:

1. Click Multicast, MVR.
2. Select Associate Profile from the Step list.
3. Select Show from the Action list.

Figure 382: Showing the MVR Group Address Profiles Assigned to a Domain

Multicast > MVR

Step: 3. Associate Profile Action: Show

Domain ID: 1

Domain Associated Profile List Total: 2

	Profile Name	Start IP Address	End IP Address
☐	rd	234.5.6.7	234.5.6.7
☐	sales	234.5.6.8	234.5.6.10

Delete Revert

Configuring MVR Interface Status Use the Multicast > MVR (Configure Interface) page to configure each interface that participates in the MVR protocol as a source port or receiver port. If you are sure that only one subscriber attached to an interface is receiving multicast services, you can enable the immediate leave function.

Command Usage

- ◆ A port configured as an MVR receiver or source port can join or leave multicast groups configured under MVR. However, note that these ports can also use IGMP snooping to join or leave any other multicast groups using the standard rules for multicast filtering.
- ◆ Receiver ports can belong to different VLANs, but should not be configured as a member of the MVR VLAN. MVR allows a receiver port to dynamically join or leave multicast groups sourced through the MVR VLAN. Multicast groups can also be statically assigned to a receiver port (see [“Assigning Static MVR Multicast Groups to Interfaces” on page 589](#)).

Receiver ports should not be statically configured as a member of the MVR VLAN. If so configured, its MVR status will be inactive. Also, note that VLAN membership for MVR receiver ports cannot be set to access mode (see [“Adding Static Members to VLANs” on page 165](#)).

- ◆ One or more interfaces may be configured as MVR source ports. A source port is able to both receive and send data for configured MVR groups or for groups which have been statically assigned (see [“Assigning Static MVR Multicast Groups to Interfaces” on page 589](#)).

All source ports must belong to the MVR VLAN.

Subscribers should not be directly connected to source ports.

- ◆ Immediate leave applies only to receiver ports. When enabled, the receiver port is immediately removed from the multicast group identified in the leave message. When immediate leave is disabled, the switch follows the standard rules by sending a query message to the receiver port and waiting for a response to determine if there are any remaining subscribers for that multicast group before removing the port from the group list.
 - Using immediate leave can speed up leave latency, but should only be enabled on a port attached to one multicast subscriber to avoid disrupting services to other group members attached to the same interface.
 - Immediate leave does not apply to multicast groups which have been statically assigned to a port.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Port/Trunk** – Interface identifier.

- ◆ **Type** – The following interface types are supported:
 - **Source** – An uplink port that can send and receive multicast data for the groups assigned to the MVR VLAN. Note that the source port must be manually configured as a member of the MVR VLAN (see [“Adding Static Members to VLANs” on page 165](#)).
 - **Receiver** – A subscriber port that can receive multicast data sent through the MVR VLAN. Any port configured as a receiver port will be dynamically added to the MVR VLAN when it forwards an IGMP report or join message from an attached host requesting any of the designated multicast services supported by the MVR VLAN. Just remember that only IGMP version 2 or 3 hosts can issue multicast join or leave messages. If MVR must be configured for an IGMP version 1 host, the multicast groups must be statically assigned (see [“Assigning Static MVR Multicast Groups to Interfaces” on page 589](#)).
 - **Non-MVR** – An interface that does not participate in the MVR VLAN. (This is the default type.)
- ◆ **Forwarding Status** – Shows if MVR traffic is being forwarded or discarded.
- ◆ **MVR Status** – Shows the MVR status. MVR status for source ports is “Active” if MVR is globally enabled on the switch. MVR status for receiver ports is “Active” only if there are subscribers receiving multicast traffic from one of the MVR groups, or a multicast group has been statically assigned to an interface.
- ◆ **Immediate Leave** – Configures the switch to immediately remove an interface from a multicast stream as soon as it receives a leave message for that group. (This option only applies to an interface configured as an MVR receiver.) (Default: Disabled)

Web Interface

To configure interface settings for MVR:

1. Click Multicast, MVR.
2. Select Configure Interface from the Step list.
3. Select an MVR domain.
4. Click Port or Trunk.
5. Set each port that will participate in the MVR protocol as a source port or receiver port, and optionally enable Immediate Leave on any receiver port to which only one subscriber is attached.
6. Click Apply.

Figure 383: Configuring Interface Settings for MVR

Multicast > MVR

Step: 5. Configure Interface

Domain ID: 1

Interface: ☒ Port ☐ Trunk

Port List Total: 28

Port	Type	Forwarding Status	MVR Status	Immediate Leave
1	Source	Forwarding	Inactive	☒ Enabled
2	Receiver	Forwarding	Inactive	☒ Enabled
3	Non-MVR	Discarding	Inactive	☒ Enabled
4	Non-MVR	Discarding	Inactive	☒ Enabled
5	Non-MVR	Discarding	Inactive	☒ Enabled

Assigning Static MVR Multicast Groups to Interfaces

Use the Multicast > MVR (Configure Static Group Member) page to statically bind multicast groups to a port which will receive long-term multicast streams associated with a stable set of hosts.

Command Usage

- ◆ Multicast groups can be statically assigned to a receiver port using this configuration page.
- ◆ The IP address range from 224.0.0.0 to 239.255.255.255 is used for multicast streams. MVR group addresses cannot fall within the reserved IP multicast address range of 224.0.0.x.
- ◆ Only IGMP version 2 or 3 hosts can issue multicast join or leave messages. If MVR must be configured for an IGMP version 1 host, the multicast groups must be statically assigned.
- ◆ The MVR VLAN cannot be specified as the receiver VLAN for static bindings.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Interface** – Port or trunk identifier.
- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Group IP Address** – Defines a multicast service sent to the selected port. Multicast groups must be assigned from the MVR group range configured on the Configure General page.

Web Interface

To assign a static MVR group to an interface:

1. Click Multicast, MVR.
2. Select Configure Static Group Member from the Step list.
3. Select Add from the Action list.
4. Select an MVR domain.
5. Select a VLAN and interface to receive the multicast stream, and then enter the multicast group address.
6. Click Apply.

Figure 384: Assigning Static MVR Groups to an Interface

The screenshot shows a web interface for configuring Multicast VLAN Registration (MVR). The breadcrumb trail is "Multicast > MVR". The "Step" dropdown is set to "5. Configure Static Group Member" and the "Action" dropdown is set to "Add". The configuration fields are as follows:

Field	Value
Domain ID	1
Interface	Port 2
VLAN	2
Group IP Address	234.5.6.8

At the bottom right, there are "Apply" and "Revert" buttons.

To show the static MVR groups assigned to an interface:

1. Click Multicast, MVR.
2. Select Configure Static Group Member from the Step list.
3. Select Show from the Action list.
4. Select an MVR domain.
5. Select the port or trunk for which to display this information.

Figure 385: Showing the Static MVR Groups Assigned to a Port

The screenshot shows the 'Multicast > MVR' configuration page. The 'Step' is set to '5. Configure Static Group Member' and the 'Action' is 'Show'. The 'Domain ID' is set to '1' and the 'Interface' is set to 'Port 2'. Below this, the 'MVR Static Group Member List' is displayed with a total of 3 members. The list has three columns: a checkbox, 'VLAN', and 'Group IP Address'. The entries are for VLAN 2 with IP addresses 234.5.6.8, 234.5.6.9, and 234.5.6.10. At the bottom right are 'Delete' and 'Revert' buttons.

	VLAN	Group IP Address
☐	2	234.5.6.8
☐	2	234.5.6.9
☐	2	234.5.6.10

Displaying MVR Receiver Groups Use the Multicast > MVR (Show Member) page to show the multicast groups either statically or dynamically assigned to the MVR receiver groups on each interface.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Group IP Address** – Multicast groups assigned to the MVR VLAN.
- ◆ **VLAN** – The VLAN through which the service is received. Note that this may be different from the MVR VLAN if the group address has been statically assigned.
- ◆ **Port** – Shows the interfaces with subscribers for multicast services provided through the MVR VLAN.
- ◆ **Up Time** – Time this service has been forwarded to attached clients.
- ◆ **Expire** – Time before this entry expires if no membership report is received from currently active or new clients.
- ◆ **Count** – The number of multicast services currently being forwarded from the MVR VLAN.
- ◆ **Clear MVR Group** – Clears multicast group information dynamically learned through MVR6. Statically configured multicast addresses are not cleared.

Web Interface

To display the interfaces assigned to the MVR receiver groups:

1. Click Multicast, MVR.
2. Select Show Member from the Step list.
3. Select an MVR domain.

Figure 386: Displaying MVR Receiver Groups

Multicast > MVR

Step: 6. Show Member

Domain ID: 1

MVR Member List Total: 8

Group IP Address	VLAN	Port	Up Time	Expire	Count
224.1.1.1	2		00:00:30		2 (Port)
	2	Unit 1 / Port 1 (Source)			
224.1.1.2	1	Unit 1 / Port 2 (Receiver)	00:01:10	00:00	4 (Host)
	4		00:00:50		4 (Port)
	4	Unit 1 / Port 3 (Source)			
	5	Unit 1 / Port 4 (Receiver)			
	6	Unit 1 / Port 5 (Source)			
	7	Unit 1 / Port 6 (Receiver)	00:01:10	00:00	1 (Host)

Displaying MVR Statistics Use the Multicast > MVR > Show Statistics pages to display MVR protocol-related statistics for the specified interface.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Port** – Port identifier. (Range: 1-28)
- ◆ **Trunk** – Trunk identifier. (Range: 1-16)

Query Statistics

- ◆ **Querier IP Address** – The IP address of the querier on this interface.
- ◆ **Querier Expire Time** – The time after which this querier is assumed to have expired.
- ◆ **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.
- ◆ **Number of Reports Sent** – The number of reports sent from this interface.
- ◆ **Number of Leaves Sent** – The number of leaves sent from 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 MVR group report received.
- ◆ **Join Success** – The number of times a multicast group was successfully joined.
- ◆ **Group** – The number of MVR 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 MVR query-related messages:

1. Click Multicast, MVR.
2. Select Show Statistics from the Step list.
3. Select Show Query Statistics from the Action list.
4. Select an MVR domain.

Figure 387: Displaying MVR Statistics – Query

The screenshot displays the 'Multicast > MVR' web interface. At the top, there are two dropdown menus: 'Step' set to '7. Show Statistics' and 'Action' set to 'Show Query Statistics'. Below these, a 'Domain ID' dropdown is set to '1'. The main section is titled 'Query Statistics' and contains a table with the following data:

Querier IP Address	None
Querier Expire Time	00(h):00(m):00(s)
General Query Received	0
General Query Sent	0
Specific Query Received	0
Specific Query Sent	0
Number of Reports Sent	0
Number of Leaves Sent	0

At the bottom left, there is a 'Clear All' button and a link that says 'Click this button to clear all MVR statistics of the domain.' At the bottom right, there is a 'Refresh' button.

To display MVR protocol-related statistics for a VLAN:

1. Click Multicast, MVR.
2. Select Show Statistics from the Step list.
3. Select Show VLAN Statistics from the Action list.
4. Select an MVR domain.
5. Select a VLAN.

Figure 388: Displaying MVR Statistics – VLAN

The screenshot shows a web interface for Multicast > MVR configuration. At the top, there's a breadcrumb 'Multicast > MVR'. Below it, there are two dropdown menus: 'Step: 7. Show Statistics' and 'Action: Show VLAN Statistics'. Underneath, there are two more dropdown menus: 'Domain ID' and 'VLAN', both set to '1'. The main content area is divided into two sections: 'Input Statistics' and 'Output Statistics'. Each section contains a table of statistics with values of 0. A 'Refresh' button is located at the bottom right.

Multicast > MVR			
Step: 7. Show Statistics		Action: Show VLAN Statistics	
Domain ID: 1		VLAN: 1	
Input Statistics			
Report	0	Drop	0
Leave	0	Join Success	0
G Query	0	Group	0
G(-S)-S Query	0		
Output Statistics			
Report	0		
Leave	0		
G Query	0		
G(-S)-S Query	0		
Refresh			

To display MVR protocol-related statistics for a port:

1. Click Multicast, MVR.
2. Select Show Statistics from the Step list.
3. Select Show Port Statistics from the Action list.
4. Select an MVR domain.
5. Select a Port.

Figure 389: Displaying MVR Statistics – Port

The screenshot shows the 'Multicast > MVR' configuration page. At the top, there are two dropdown menus: 'Step: 7. Show Statistics' and 'Action: Show Port Statistics'. Below these, there are two more dropdown menus: 'Domain ID: 1' and 'Port: 1'. The main content area is divided into two sections: 'Input Statistics' and 'Output Statistics'. Each section contains a table of statistics.

Input Statistics	
Report	0
Leave	0
G Query	0
G(-S)-S Query	0
Drop	0
Join Success	0
Group	0

Output Statistics	
Report	0
Leave	0
G Query	0
G(-S)-S Query	0

At the bottom right of the statistics section, there is a 'Refresh' button.

Multicast VLAN Registration for IPv6

MVR6 functions in a manner similar to that described for MVR (see [“Multicast VLAN Registration for IPv4” on page 579](#)).

Command Usage

◆ General Configuration Guidelines for MVR6:

1. Enable MVR6 for a domain on the switch, and select the MVR VLAN (see [“Configuring MVR6 Domain Settings” on page 599](#)).
2. Create an MVR6 profile by specifying the multicast groups that will stream traffic to attached hosts, and assign the profile to an MVR6 domain (see [“Configuring MVR6 Group Address Profiles” on page 600](#)).

3. Set the interfaces that will join the MVR as source ports or receiver ports (see [“Configuring MVR6 Interface Status” on page 603](#)).
4. For multicast streams that will run for a long term and be associated with a stable set of hosts, you can statically bind the multicast group to the participating interfaces (see [“Assigning Static MVR6 Multicast Groups to Interfaces” on page 605](#)).

Configuring MVR6 Global Settings Use the Multicast > MVR6 (Configure Global) page to configure proxy switching and the robustness variable.

Parameters

These parameters are displayed:

- ◆ **Proxy Switching** – Configures MVR proxy switching, where the source port acts as a host, and the receiver port acts as an MVR router with querier service enabled. (Default: Enabled)
 - When MVR proxy-switching is enabled, an MVR source port serves as the upstream or host interface, and the MVR receiver port serves as the querier. The source port performs only the host portion of MVR by sending summarized membership reports, and automatically disables MVR router functions.
 - Receiver ports are known as downstream or router interfaces. These interfaces perform the standard MVR router functions by maintaining a database of all MVR subscriptions on the downstream interface. Receiver ports must therefore be configured on all downstream interfaces which require MVR proxy service.
 - When the source port receives report and leave messages, it only forwards them to other source ports.
 - When receiver ports receive any query messages, they are dropped.
 - When changes occurring in the downstream MVR groups are learned by the receiver ports through report and leave messages, an MVR state change report is created and sent to the upstream source port, which in turn forwards this information upstream.
 - When MVR proxy switching is disabled:
 - Any membership reports received from receiver/source ports are forwarded to all source ports.
 - When a source port receives a query message, it will be forwarded to all downstream receiver ports.
 - When a receiver port receives a query message, it will be dropped.

- ◆ **Robustness Value** – Configures the expected packet loss, and thereby the number of times to generate report and group-specific queries. (Range: 1-10; Default: 2)
 - This parameter is used to set the number of times report messages are sent upstream when changes are learned about downstream groups, and the number of times group-specific queries are sent to downstream receiver ports.
 - This parameter only takes effect when MVR6 proxy switching is enabled.
- ◆ **Proxy Query Interval** – Configures the interval at which the receiver port sends out general queries. (Range: 2-31744 seconds; Default: 125 seconds)
 - This parameter sets the general query interval at which active receiver ports send out general queries.
 - This interval is only effective when proxy switching is enabled.
- ◆ **Source Port Mode** – Configures the switch to forward any multicast streams within the parameters set by a profile, or to only forward multicast streams which the source port has dynamically joined.
 - **Always Forward** – By default, the switch forwards any multicast streams within the address range set by a profile, and bound to a domain. The multicast streams are sent to all source ports on the switch and to all receiver ports that have elected to receive data on that multicast address. (This is the default setting.)
 - **Dynamic** – When dynamic mode is enabled, the switch only forwards multicast streams which the source port has dynamically joined. In other words, both the receiver port and source port must subscribe to a multicast group before a multicast stream is forwarded to any attached client. Note that the requested streams are still restricted to the address range which has been specified in a profile and bound to a domain.

Web Interface

To configure global settings for MVR6:

1. Click Multicast, MVR6.
2. Select Configure Global from the Step list.
3. Set the status for MVR6 proxy switching, the robustness value used for report and query messages, the proxy query interval, and source port mode.
4. Click Apply.

Figure 390: Configuring Global Settings for MVR6

The screenshot shows the 'Multicast > MVR6' configuration page. At the top, there is a breadcrumb 'Multicast > MVR6'. Below it is a 'Step:' dropdown menu set to '1. Configure Global'. The main configuration area contains the following settings:

- Proxy Switching:** A checkbox labeled 'Enabled' is checked.
- Robustness Value (1-10):** A text input field containing the value '2'.
- Proxy Query Interval (2-31744):** A text input field containing the value '125', followed by a 'sec' label.
- Source Port Mode:** A dropdown menu set to 'Always Forward'.

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

Configuring MVR6 Domain Settings

Use the Multicast > MVR6 (Configure Domain) page to enable MVR6 globally on the switch, and select the VLAN that will serve as the sole channel for common multicast streams supported by the service provider.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **MVR6 Status** – When MVR6 is enabled on the switch, any multicast data associated with an MVR6 group is sent from all designated source ports, to all receiver ports that have registered to receive data from that multicast group. (Default: Disabled)
- ◆ **MVR6 VLAN** – Identifier of the VLAN that serves as the channel for streaming multicast services using MVR6. MVR6 source ports should be configured as members of the MVR6 VLAN (see [“Adding Static Members to VLANs” on page 165](#)), but MVR6 receiver ports should not be manually configured as members of this VLAN. (Default: 1)
- ◆ **MVR6 Running Status** – Indicates whether or not all necessary conditions in the MVR6 environment are satisfied. Running status is Active as long as MVR6 is enabled, the specified MVR6 VLAN exists, and a source port with a valid link has been configured (see [“Configuring MVR6 Interface Status” on page 603](#)).
- ◆ **MVR6 Current Learned Groups** – The number of MVR6 groups currently assigned to this domain.
- ◆ **Forwarding Priority** – The CoS priority assigned to all multicast traffic forwarded into this domain. (Range: 0-7, where 7 is the highest priority)
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.
- ◆ **Upstream Source IPv6** – The source IPv6 address assigned to all MVR6 control packets sent upstream on the specified domain. This parameter must be a full

IPv6 address including the network prefix and host address bits. By default, all MVR6 reports sent upstream use a null source IP address.

All IPv6 addresses must be 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. (Note that the IP address ff02::X is reserved.)

Web Interface

To configure settings for an MVR6 domain:

1. Click Multicast, MVR6.
2. Select Configure Domain from the Step list.
3. Select a domain from the scroll-down list.
4. Enable MVR6 for the selected domain, select the MVR6 VLAN, set the forwarding priority to be assigned to all ingress multicast traffic, and set the source IP address for all control packets sent upstream as required.
5. Click Apply.

Figure 391: Configuring Domain Settings for MVR6

The screenshot shows the 'Multicast > MVR6' configuration page. The 'Step' dropdown is set to '2. Configure Domain'. The configuration fields are as follows:

Field	Value
Domain ID	1
MVR6 Status	☒ Enabled
MVR6 VLAN	1
MVR6 Running Status	Inactive
MVR6 Current Learned Groups	0
Forwarding Priority (0-7)	☐ 0
Upstream Source IPv6	2001:d88:2222:7223::72

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

Configuring MVR6 Group Address Profiles

Use the Multicast > MVR6 (Configure Profile and Associate Profile) pages to assign the multicast group address for required services to one or more MVR6 domains.

Command Usage

- ◆ Use the Configure Profile page to statically configure all multicast group addresses that will join the MVR6 VLAN. Any multicast data associated with an MVR6 group is sent from all source ports to all receiver ports that have registered to receive data from that multicast group.

- ◆ All IPv6 addresses must be 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. (Note that the IP address ff02::X is reserved.)
- ◆ The MVR6 group address range assigned to a profile cannot overlap with the group address range of any other profile.
- ◆ MVR6 domains can be associated with more than one MVR6 profile. But since MVR6 domains cannot share the group range, an MVR6 profile can only be associated with one MVR6 domain.

Parameters

These parameters are displayed:

Configure Profile

- ◆ **Profile Name** – The name of a profile containing one or more MVR6 group addresses. (Range: 1-21 characters)
- ◆ **Start IPv6 Address** – Starting IP address for an MVR6 multicast group. This parameter must be a full IPv6 address including the network prefix and host address bits.
- ◆ **End IPv6 Address** – Ending IP address for an MVR6 multicast group. This parameter must be a full IPv6 address including the network prefix and host address bits.

Associate Profile

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Profile Name** – The name of a profile to be assigned to this domain. (Range: 1-21 characters)

Web Interface

To configure an MVR6 group address profile:

1. Click Multicast, MVR6.
2. Select Configure Profile from the Step list.
3. Select Add from the Action list.
4. Enter the name of a group profile to be assigned to one or more domains, and specify a multicast group that will stream traffic to participating hosts.
5. Click Apply.

Figure 392: Configuring an MVR6 Group Address Profile

Multicast > MVR6

Step: 2. Configure Profile Action: Add

Profile Name: rd

Start IPv6 Address: ff00::1

End IPv6 Address: ff00::9

Apply Revert

To show the configured MVR6 group address profiles:

1. Click Multicast, MVR6.
2. Select Configure Profile from the Step list.
3. Select Show from the Action list.

Figure 393: Displaying MVR6 Group Address Profiles

Multicast > MVR6

Step: 3. Configure Profile Action: Show

MVR6 Profile List Total: 1

☐	Profile Name	Start IPv6 Address	End IPv6 Address
☐	rd	FF00::1	FF00::9

Delete Revert

To assign an MVR6 group address profile to a domain:

1. Click Multicast, MVR6.
2. Select Associate Profile from the Step list.
3. Select Add from the Action list.
4. Select a domain from the scroll-down list, and enter the name of a group profile.
5. Click Apply.

Figure 394: Assigning an MVR6 Group Address Profile to a Domain

Multicast > MVR6

Step: 4. Associate Profile Action: Add

Domain ID: 1

Profile Name: rd

Apply Revert

To show the MVR6 group address profiles assigned to a domain:

1. Click Multicast, MVR6.
2. Select Associate Profile from the Step list.
3. Select Show from the Action list.

Figure 395: Showing MVR6 Group Address Profiles Assigned to a Domain

Multicast > MVR6

Step: 3. Associate Profile Action: Show

Domain ID: 1

Domain Associated Profile List Total: 1

Profile Name	Start IPv6 Address	End IPv6 Address
rd	FF00::1	FF00::9

Delete Revert

Configuring MVR6 Interface Status

Use the Multicast > MVR6 (Configure Interface) page to configure each interface that participates in the MVR6 protocol as a source port or receiver port. If you are sure that only one subscriber attached to an interface is receiving multicast services, you can enable the immediate leave function.

Command Usage

- ◆ A port configured as an MVR6 receiver or source port can join or leave multicast groups configured under MVR6. A port which is not configured as an MVR receiver or source port can use MLD snooping to join or leave multicast groups using the standard rules for multicast filtering (see [“MLD Snooping \(Snooping and Query for IPv6\)”](#) on page 561).
- ◆ Receiver ports can belong to different VLANs, but should not be configured as a member of the MVR6 VLAN. MVR6 allows a receiver port to dynamically join or leave multicast groups sourced through the MVR6 VLAN. Multicast groups can also be statically assigned to a receiver port (see [“Assigning Static MVR Multicast Groups to Interfaces”](#) on page 589).

Receiver ports should not be statically configured as a member of the MVR6 VLAN. If so configured, its MVR6 status will be inactive. Also, note that VLAN

membership for MVR6 receiver ports cannot be set to access mode (see [“Adding Static Members to VLANs” on page 165](#)).

- ◆ One or more interfaces may be configured as MVR6 source ports. A source port is able to both receive and send data for configured MVR6 groups or for groups which have been statically assigned (see [“Assigning Static MVR Multicast Groups to Interfaces” on page 589](#)).

All source ports must belong to the MVR6 VLAN.

Subscribers should not be directly connected to source ports.

- ◆ Immediate leave applies only to receiver ports. When enabled, the receiver port is immediately removed from the multicast group identified in the leave message. When immediate leave is disabled, the switch follows the standard rules by sending a group-specific query to the receiver port and waiting for a response to determine if there are any remaining subscribers for that multicast group before removing the port from the group list.
 - Using immediate leave can speed up leave latency, but should only be enabled on a port attached to one multicast subscriber to avoid disrupting services to other group members attached to the same interface.
 - Immediate leave does not apply to multicast groups which have been statically assigned to a port.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Port/Trunk** – Interface identifier.
- ◆ **Type** – The following interface types are supported:
 - **Non-MVR6** – An interface that does not participate in the MVR6 VLAN. (This is the default type.)
 - **Source** – An uplink port that can send and receive multicast data for the groups assigned to the MVR6 VLAN. Note that the source port must be manually configured as a member of the MVR6 VLAN (see [“Adding Static Members to VLANs” on page 165](#)).
 - **Receiver** – A subscriber port that can receive multicast data sent through the MVR6 VLAN. Also, note that VLAN membership for MVR receiver ports cannot be set to access mode (see [“Adding Static Members to VLANs” on page 165](#)).
- ◆ **Forwarding Status** – Shows if multicast traffic is being forwarded or blocked.
- ◆ **MVR6 Status** – Shows the MVR6 status. MVR6 status for source ports is “Active” if MVR6 is globally enabled on the switch. MVR6 status for receiver ports is

“Active” only if there are subscribers receiving multicast traffic from one of the MVR6 groups, or a multicast group has been statically assigned to an interface.

- ◆ **Immediate Leave** – Configures the switch to immediately remove an interface from a multicast stream as soon as it receives a leave message for that group. This option only applies to an interface configured as an MVR6 receiver. (Default: Disabled)

Web Interface

To configure interface settings for MVR6:

1. Click Multicast, MVR6.
2. Select Configure Interface from the Step list.
3. Select an MVR6 domain.
4. Click Port or Trunk.
5. Set each port that will participate in the MVR6 protocol as a source port or receiver port, and optionally enable Immediate Leave on any receiver port to which only one subscriber is attached.
6. Click Apply.

Figure 396: Configuring Interface Settings for MVR6

Multicast > MVR6

Step: 5. Configure Interface

Domain ID: 1

Interface: ☒ Port ☐ Trunk

Port List Total: 28

Port	Type	Forwarding Status	MVR6 Status	Immediate Leave
1	Source	Forwarding	Inactive	☐ By-Group
2	Receiver	Forwarding	Inactive	☒ By-Group
3	Non-MVR6	Discarding	Inactive	☐ By-Group
4	Non-MVR6	Discarding	Inactive	☐ By-Group
5	Non-MVR6	Discarding	Inactive	☐ By-Group

Assigning Static MVR6 Multicast Groups to Interfaces

Use the Multicast > MVR6 (Configure Static Group Member) page to statically bind multicast groups to a port which will receive long-term multicast streams associated with a stable set of hosts.

Command Usage

- ◆ Multicast groups can be statically assigned to a receiver port using this configuration page.

- ◆ All IPv6 addresses must be 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. (Note that the IP address ff02::X is reserved.)
- ◆ The MVR6 VLAN cannot be specified as the receiver VLAN for static bindings.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Interface** – Port or trunk identifier.
- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Group IPv6 Address** – Defines a multicast service sent to the selected port. Multicast groups must be assigned from the MVR6 group range configured on the Configure General page.

Web Interface

To assign a static MVR6 group to an interface:

1. Click Multicast, MVR6.
2. Select Configure Static Group Member from the Step list.
3. Select Add from the Action list.
4. Select an MVR6 domain.
5. Select a VLAN and interface to receive the multicast stream, and then enter the multicast group address.
6. Click Apply.

Figure 397: Assigning Static MVR6 Groups to a Port

The screenshot shows a web interface for configuring MVR6. The breadcrumb is 'Multicast > MVR6'. Below it, there's a 'Step:' dropdown set to '5. Configure Static Group Member' and an 'Action:' dropdown set to 'Add'. The configuration fields are as follows:

Domain ID	1
Interface	☒ Port 2 ☐ Trunk
VLAN	2
Group IPv6 Address	ff00::1

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

To show the static MVR6 groups assigned to an interface:

1. Click Multicast, MVR6.
2. Select Configure Static Group Member from the Step list.
3. Select Show from the Action list.
4. Select an MVR6 domain.
5. Select the port or trunk for which to display this information.

Figure 398: Showing the Static MVR6 Groups Assigned to a Port

Multicast > MVR6

Step: 6. Configure Static Group Member Action: Show

Domain ID: 1

Interface: Port 2

MVR6 Static Group Member List Total: 1

	VLAN	Group IPv6 Address
☐	2	FF00::1

Delete Revert

Displaying MVR6 Receiver Groups

Use the Multicast > MVR6 (Show Member) page to show the multicast groups either statically or dynamically assigned to the MVR6 receiver groups on each interface.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Group IPv6 Address** – Multicast groups assigned to the MVR6 VLAN.
- ◆ **VLAN** – The VLAN through which the service is received. Note that this may be different from the MVR6 VLAN if the group address has been statically assigned.
- ◆ **Port** – Indicates the source address of the multicast service, or displays an asterisk if the group address has been statically assigned (these entries are marked as “Source”). Also shows the interfaces with subscribers for multicast services provided through the MVR6 VLAN (these entries are marked as “Receiver”).
- ◆ **Up Time** – Time this service has been forwarded to attached clients.
- ◆ **Expire** – Time before this entry expires if no membership report is received from currently active or new clients.

- ◆ **Count** – The number of multicast services currently being forwarded from the MVR6 VLAN.
- ◆ **Clear MVR6 Group** – Clears multicast group information dynamically learned through MVR6. Statically configured multicast addresses are not cleared.

Web Interface

To display the interfaces assigned to the MVR6 receiver groups:

1. Click Multicast, MVR6.
2. Select Show Member from the Step list.
3. Select an MVR6 domain.

Figure 399: Displaying MVR6 Receiver Groups

Multicast > MVR6					
Step: 6. Show Member					
Domain ID: 1					
MVR6 Member List Total: 3					
Group IPv6 Address	VLAN	Port	Up Time	Expire	Count
FF00::1	1	Unit 1 / Port 1 (Source)	00:00:11:35		2 (Port)
	2	Unit 1 / Port 2 (Receiver)			0 (Host)

Displaying MVR6 Statistics Use the Multicast > MVR6 > Show Statistics pages to display MVR6 protocol-related statistics for the specified interface.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Port** – Port identifier. (Range: 1-28)
- ◆ **Trunk** – Trunk identifier. (Range: 1-16)

Query Statistics

- ◆ **Querier IPv6 Address** – The IP address of the querier on this interface.
- ◆ **Querier Expire Time** – The time after which this querier is assumed to have expired.
- ◆ **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.
- ◆ **Number of Reports Sent** – The number of reports sent from this interface.
- ◆ **Number of Leaves Sent** – The number of leaves sent from this interface.

VLAN, Port, and Trunk Statistics

Input Statistics

- ◆ **Report** – The number of MLD 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 MVR6 group report received.
- ◆ **Join Success** – The number of times a multicast group was successfully joined.
- ◆ **Group** – The number of MVR6 groups active on this interface.

Output Statistics

- ◆ **Report** – The number of MLD 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 MVR6 query-related messages:

1. Click Multicast, MVR6.
2. Select Show Statistics from the Step list.
3. Select Show Query Statistics from the Action list.
4. Select an MVR6 domain.

Figure 400: Displaying MVR6 Statistics – Query

The screenshot displays the 'Multicast > MVR6' web interface. At the top, there are two dropdown menus: 'Step' set to '7. Show Statistics' and 'Action' set to 'Show Query Statistics'. Below these, a 'Domain ID' dropdown is set to '1'. The main section is titled 'Query Statistics' and contains a table with the following data:

Querier IPv6 Address	::
Querier Expire Time	00(h):00(m):00(s)
General Query Received	0
General Query Sent	34
Specific Query Received	0
Specific Query Sent	0
Number of Reports Sent	0
Number of Leaves Sent	0

At the bottom left, there is a 'Clear All' button and a link that says 'Click this button to clear all MVR6 statistics of the domain.' At the bottom right, there is a 'Refresh' button.

To display MVR6 protocol-related statistics for a VLAN:

1. Click Multicast, MVR6.
2. Select Show Statistics from the Step list.
3. Select Show VLAN Statistics from the Action list.
4. Select an MVR6 domain.
5. Select a VLAN.

Figure 401: Displaying MVR6 Statistics – VLAN

The screenshot shows the 'Multicast > MVR6' configuration page. At the top, there are two dropdown menus: 'Step: 7. Show Statistics' and 'Action: Show VLAN Statistics'. Below these, there are two more dropdown menus: 'Domain ID' and 'VLAN', both set to '1'. The main content area displays statistics for the selected domain and VLAN. It is divided into two sections: 'Input Statistics' and 'Output Statistics'. The 'Input Statistics' section shows values for Report (0), Done (0), G Query (0), G(-S)-S Query (0), Drop (0), Join Success (0), and Group (1). The 'Output Statistics' section shows values for Report (0), Done (0), G Query (35), and G(-S)-S Query (0). At the bottom right, there are two buttons: 'Clear' and 'Refresh'.

Multicast > MVR6			
Step: 7. Show Statistics		Action: Show VLAN Statistics	
Domain ID: 1		VLAN: 1	
Input Statistics			
Report	0	Drop	0
Done	0	Join Success	0
G Query	0	Group	1
G(-S)-S Query	0		
Output Statistics			
Report	0		
Done	0		
G Query	35		
G(-S)-S Query	0		

Clear Refresh

To display MVR6 protocol-related statistics for a port:

1. Click Multicast, MVR6.
2. Select Show Statistics from the Step list.
3. Select Show Port Statistics from the Action list.
4. Select an MVR6 domain.
5. Select a Port.

Figure 402: Displaying MVR6 Statistics – Port

The screenshot shows the 'Multicast > MVR6' configuration page. At the top, there are two dropdown menus: 'Step: 7. Show Statistics' and 'Action: Show Port Statistics'. Below these, there are two more dropdown menus: 'Domain ID: 1' and 'Port: 1'. The main content area is divided into two sections: 'Input Statistics' and 'Output Statistics'. The 'Input Statistics' section shows four rows of data: 'Report' (0), 'Done' (0), 'G Query' (0), and 'G(-S)-S Query' (0). The 'Output Statistics' section shows four rows of data: 'Report' (12), 'Done' (1), 'G Query' (0), and 'G(-S)-S Query' (0). At the bottom right, there are two buttons: 'Clear' and 'Refresh'.

Multicast > MVR6			
Step: 7. Show Statistics		Action: Show Port Statistics	
Domain ID: 1		Port: 1	
Input Statistics			
Report	0	Drop	0
Done	0	Join Success	0
G Query	0	Group	0
G(-S)-S Query	0		
Output Statistics			
Report	12		
Done	1		
G Query	0		
G(-S)-S Query	0		

Clear Refresh

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 BOOTP or DHCP server. An IPv6 address can only be manually configured.

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 618](#).

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 BOOTP or 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 675](#)) or use dynamic routing; i.e., RIP ([page 692](#)), OSPFv2 ([page 710](#)), 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 Address) menu, static routes ([page 675](#)), 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), 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 403: Configuring a Static IPv4 Address

The screenshot shows the 'IP > General > Routing Interface' configuration page. The 'Action' dropdown is set to 'Add Address'. The 'VLAN' dropdown is set to '1'. The 'IP Address Mode' dropdown is set to 'User Specified'. The 'IP Address Type' dropdown is set to 'Primary'. The 'IP Address' text field contains '192.168.0.2'. The 'Subnet Mask' text field contains '255.255.255.0'. Below these fields is a 'Restart DHCP' button and a link that says 'Click this button to restart DHCP service.' At the bottom right are 'Apply' and 'Revert' buttons.

To obtain an 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 404: Configuring a Dynamic IPv4 Address

IP > General > Routing Interface

Action: Add Address

VLAN: 1

IP Address Mode: DHCP

IP Address Type: Primary

IP Address:

Subnet Mask:

Restart DHCP [Click this button to restart DHCP service.](#)

Apply Revert



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 405: Showing the Configured IPv4 Address for an Interface

IP > General > Routing Interface

Action:

VLAN:

IP Address Mode: User Specified

Routing Interface IP List Total: 1

	IP Address Type	IP Address	Subnet Mask
☐	Primary	192.168.0.2	255.255.255.0

Sending DHCP Inform Requests for Additional Information

Use the IP > General > Routing Interface (Configure Interface) page to submit a DHCP request for information about the default domain name server and default gateway from a VLAN interface configured with a static IPv4 address.

Command Usage

- ◆ A DHCP Inform request can only be issued for a VLAN interface which has been configured with a static IPv4 address.
- ◆ The switch will continue sending DHCP inform requests until it receives an acknowledgement from the server. If the retry interval is larger than 60 seconds, the switch will stop sending requests, record a log entry, and then wait for another 300 seconds before trying again.
- ◆ When the server receives a DHCP inform request, it will return the requested optional information to the host without assigning an IP address. If there is no option 6 (DNS) in the DHCP acknowledgement, the switch will delete all configured name servers. If there is more than one gateway address in option 3 (default gateway), the switch will use the first entry.
- ◆ If the interface address is changed to another static IP address in a different subnet, the DHCP inform request should be sent again.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of configured VLAN (1-4094).
- ◆ **DHCP Inform** – Submits a DHCP request for information about the default domain name server and default gateway from a VLAN interface configured with a static IPv4 address. (Default: Disabled)

Web Interface

To send DHCP Inform requests for additional information:

1. Click IP, General, Routing Interface.
2. Select Configure Interface from the Action list.
3. Select a VLAN configured with a static IPv4 address.
4. Set the DHCP inform field to the required status.
5. Click Apply.

Figure 406: Send DHCP Inform Requests for Additional Information

IP > General > Routing Interface

Action: Configure Interface

VLAN 1

DHCP Inform ☒ Enabled

Apply Revert

Setting the Switch's IP Address (IP Version 6)

This section describes how to configure an 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 613](#).

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 be manually configured on this switch (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 routes defined, you must define a gateway if the target device is located in a different subnet.
 - If a routing protocol is enabled ([page 691](#)), you can still define a static route ([page 675](#)) 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 407: Configuring the IPv6 Default Gateway

The screenshot shows a web interface for configuring IPv6. At the top, it says "IP > IPv6". Below that, there is a section titled "Action:" with a dropdown menu set to "Configure Global". Underneath, there is a field labeled "Default Gateway" with the value "2001:DB8:2222:7272::254" entered. At the bottom right, 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 624](#)).
- ◆ 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)
- ◆ **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, 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 162](#)). 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 408: Configuring General Settings for an IPv6 Interface

IP > IPv6 Configuration

Action: Configure Interface

Mode ☒ VLAN ☐ RA Guard

VLAN 1

Enable IPv6 Explicitly ☐ Enabled

MTU (1280-65535) 1500 bytes

ND DAD Attempts (0-600) 1

ND NS Interval (1000-3600000) 1000 ms

ND Reachable-Time (0-3600000) 30000 ms

Apply 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 409: Configuring RA Guard for an IPv6 Interface

IP > IPv6 Configuration

Action: Configure Interface

Mode ☐ VLAN ☒ RA Guard

Interface ☒ Port ☐ Trunk

Port List Total: 28

Port	RA Guard
1	☐ Enabled
2	☐ Enabled
3	☐ Enabled
4	☐ Enabled
5	☒ Enabled

Configuring an IPv6 Address Use the IP > IPv6 Configuration (Add IPv6 Address) page to configure an 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 be configured with a link-local address. Therefore, any configuration process that enables IPv6 functionality, including explicitly enabling IPv6 (see ["Configuring IPv6 Interface Settings" on page 620](#)), 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:
 - The global unicast address can be automatically configured by taking the network prefix from router advertisements observed on the local interface, and using the modified EUI-64 form of the interface identifier to automatically create the host portion of the address (see ["Configuring IPv6 Interface Settings" on page 620](#)).
 - 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 410: Configuring an IPv6 Address

The screenshot shows a web interface for configuring an IPv6 address. The breadcrumb navigation at the top reads 'IP > IPv6'. Below this, there is a section titled 'Action:' with a dropdown menu set to 'Add IPv6 Address'. Underneath, there are three configuration fields: 'VLAN' with a dropdown menu set to '1', 'Address Type' with a dropdown menu set to 'Global', and 'IPv6 Address' with a text input field containing '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 node 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 “show ip interface” command described 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 411: 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	<p>The following states are used for dynamic entries:</p> <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 Reachable 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.♦ 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 412: Showing IPv6 Neighbors



The screenshot shows the 'IP > IPv6 Configuration' page. Under the 'Action:' dropdown, 'Show IPv6 Neighbor Cache' is selected. Below this, the 'Current Neighbor Cache Table' is displayed with a 'Total: 1' count. The table has five columns: IPv6 Address, Age, Link-layer Address, State, and VLAN. One entry is shown with the IPv6 address 2001:DB8:2222:7272::74, an age of 20, a link-layer address of 00-E0-0C-9C-CA-10, a state of 'Reachable', and VLAN 1. A 'Clear' button is located at the bottom right of the table.

IPv6 Address	Age	Link-layer Address	State	VLAN
2001:DB8:2222:7272::74	20	00-E0-0C-9C-CA-10	Reachable	1

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.
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.
Reassembled 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.
Reassembled 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.

Table 37: Show IPv6 Statistics - display description (Continued)

Field	Description
<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 checksums, bad length, etc.).
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.

Table 37: Show IPv6 Statistics - display description (Continued)

Field	Description
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.
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.
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.

Table 37: Show IPv6 Statistics - display description (Continued)

Field	Description
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 413: Showing IPv6 Statistics (IPv6)

The screenshot shows the 'IP > IPv6' configuration page. At the top, there is a breadcrumb 'IP > IPv6'. Below it, the 'Action:' dropdown is set to 'Show Statistics'. Under the 'Type' section, three radio buttons are present: 'IPv6' (selected), 'ICMPv6', and 'UDP'. The main content area is titled 'IPv6 Statistics' and contains a table with two columns of statistics. The first column lists the statistic name, the second column shows the count, and the third column shows the description. The statistics are as follows:

Statistic	Count	Description
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	

At the bottom right of the statistics table, there is a 'Clear' button.

Figure 414: Showing IPv6 Statistics (ICMPv6)





Figure 415: 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 416: Showing Reported MTU Values



The screenshot shows the 'IP > IPv6' configuration page. Under the 'Action:' dropdown, 'Show MTU' is selected. Below this, the 'MTU Table' is displayed with a 'Total: 2' count. The table has three columns: 'MTU', 'Since', and 'Destination Address'. It contains two rows of data.

IP > IPv6		
Action: Show MTU		
MTU Table Total: 2		
MTU	Since	Destination Address
1400	00:04:21	5000:1::3
1280	00:04:50	FE80::203:A0FF:FED8:141D

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** – Configures client, relay, and DHCP server.
- ◆ **UDP Helper** – Configures the switch to forward UDP broadcast packets originating from host applications to another part of the network.
- ◆ **PPPoE Intermediate Agent** – Configures PPPoE Intermediate Agent (PPPoE IA) relay parameters required for passing authentication messages between a client and broadband remote access servers.

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 640](#)).
- ◆ 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 417: Configuring General Settings for DNS

The screenshot shows a web interface for configuring DNS settings. At the top, the breadcrumb is 'IP Service > DNS > General'. Below this, there is a section titled 'Action:' with a dropdown menu currently showing 'Configure Global'. Underneath, there are two settings: 'Domain Lookup' which has a checked checkbox and the text 'Enabled' next to it, and 'Default Domain Name' which has a text input field containing 'my.site.com'. At the bottom right of the configuration area, there are two buttons: 'Apply' and 'Revert'.

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 637](#)). 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 640](#)).

- ◆ If all name servers are deleted, DNS will automatically be disabled.

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 418: Configuring a List of Domain Names for DNS

IP Service > DNS > General

Action: Add Domain Name

Domain Name: sample.com.uk

Apply Revert

To show the list domain names:

1. Click IP Service, DNS.
2. Select Show Domain Names from the Action list.

Figure 419: Showing the List of Domain Names for DNS

IP Service > DNS > General

Action: Show Domain Names

Domain Name List Total: 2

	Domain Name
☐	google.com
☐	hinet.net

Delete Revert

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 637](#)).
- ◆ 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 420: 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 titled '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 421: Showing the List of Name Servers for DNS

The screenshot shows the 'IP Service > DNS > General' configuration page. At the top, there is a breadcrumb trail and an 'Action:' dropdown menu set to 'Show Name Servers'. Below this, a table titled 'Name Server IP Address List' with a 'Total: 3' indicator is displayed. The table has two columns: a checkbox column and a 'Name Server IP Address' column. It contains three entries with IP addresses 192.168.1.10, 140.113.5.7, and 10.7.231.5. At the bottom right of the table area, there are 'Delete' and 'Revert' buttons.

	Name Server IP Address
☐	192.168.1.10
☐	140.113.5.7
☐	10.7.231.5

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 422: Configuring Static Entries in the DNS Table

The screenshot shows the 'IP Service > DNS > Static Host Table' configuration page. At the top, there is a breadcrumb trail and an 'Action:' dropdown menu set to 'Add'. Below this, there are two input fields: 'Host Name' with the value 'yahoo.com' and 'IP Address' with the value '10.2.78.3'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show static entries in the DNS table:

1. Click IP Service, DNS, Static Host Table.
2. Select Show from the Action list.

Figure 423: Showing Static Entries in the DNS Table



The screenshot shows a web interface for the 'Static Host Table'. At the top, there is a breadcrumb trail 'IP Service > DNS > Static Host Table' and a green help icon. Below this is an 'Action:' dropdown menu set to 'Show'. The main content area is titled 'IP Address List Total: 3'. It contains a table with three columns: a checkbox column, a 'Host' column, and an 'IP Address' column. There are three rows of data. At the bottom right of the table area are 'Delete' and 'Revert' buttons.

	Host	IP Address
☐	yahoo.com	10.2.78.3
☐	hinet.net	124.29.31.155
☐	google.com	133.45.211.18

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 424: Showing Entries in the DNS Cache



The screenshot shows a web interface titled "IP Service > DNS > Cache". Below the title bar, it says "Cache Information Total: 3". There is a table with 6 columns: No., Flag, Type, IP, TTL, and Host. The table contains 3 rows of data. Below the table is a "Clear" button.

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

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, or configure the DHCP server on this switch to support that subnet.

When configuring the DHCP server on this switch, you can configure an address pool for each unique IP interface, or manually assign a static IP address to clients based on their hardware address or client identifier. The DHCP server can provide the host's IP address, domain name, gateway router and DNS server, information about the host's boot image including the TFTP server to access for download and the name of the boot file, or boot information for NetBIOS Windows Internet Naming Service (WINS).

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. (Range: 1-64 characters)

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 425: Specifying a DHCP Client Identifier



IP Service > DHCP > Client

VLAN: 1


Vendor Class ID: ☒ Default GTL-2891

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 426: 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 DHCP servers or relay servers in order of preference for any VLAN.
3. Click Apply.

Figure 427: 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	192.168.2.33 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0

Restart DHCP Relay Click the button to restart DHCP Relay service.


Apply Revert

Configuring the DHCP Server

This switch includes a Dynamic Host Configuration Protocol (DHCP) server that can assign temporary IP addresses to any attached host requesting service. It can also provide other network settings such as the domain name, default gateway, Domain Name Servers (DNS), Windows Internet Naming Service (WINS) name servers, or information on the bootup file for the host device to download.

Addresses can be assigned to clients from a common address pool configured for a specific IP interface on this switch, or fixed addresses can be assigned to hosts based on the client identifier code or MAC address.

Figure 428: DHCP Server



Command Usage

- ◆ First configure any excluded addresses, including the address for this switch.
- ◆ Then configure address pools for the network interfaces. You can configure up to 8 network address pools. You can also manually bind an address to a specific client if required. However, any fixed addresses must fall within the range of an existing network address pool. You can configure up to 32 fixed host addresses (i.e., entering one address per pool).
- ◆ If the DHCP server is running, you must disable it and then re-enable it to implement any configuration changes. This can be done on the IP Service > DHCP > Server (Configure Global) page.

Enabling the Server

Use the IP Service > DHCP > Server (Configure Global) page to enable the DHCP Server.

Parameters

These parameters are displayed:

- ◆ **DHCP Server** – Enables or disables the DHCP server on this switch.
(Default: Disabled)

Web Interface

To enable the DHCP server:

1. Click IP Service, DHCP, Server.
2. Select Configure Global from the Step list.
3. Mark the Enabled box.
4. Click Apply.

Figure 429: Enabling the DHCP Server

IP Service > DHCP > Server

Step: 1. Configure Global

DHCP Server ☒ Enabled

Note:
You must configure an address pool before enabling the DHCP server.
If the DHCP server is running, you must restart it to implement any configuration changes.

Apply Revert

Setting Excluded Addresses

Use the IP Service > DHCP > Server (Configure Excluded Addresses – Add) page to specify the IP addresses that should not be assigned to clients.

Parameters

These parameters are displayed:

- ◆ **Start IP Address** – Specifies a single IP address or the first address in a range that the DHCP server should not assign to DHCP clients.
- ◆ **End IP Address** – The last address in a range that the DHCP server should not assign to DHCP clients.



Note: Be sure you exclude the address for this switch and other key network devices.

Web Interface

To configure IP addresses excluded for DHCP clients:

1. Click IP Service, DHCP, Server.
2. Select Configure Excluded Addresses from the Step list.
3. Select Add from the Action list.
4. Enter a single address or an address range.
5. Click Apply.

Figure 430: Configuring Excluded Addresses on the DHCP Server

The screenshot shows a web interface for configuring DHCP excluded addresses. At the top, the breadcrumb 'IP Service > DHCP > Server' is displayed. Below it, the 'Step' dropdown is set to '2. Configure Excluded Address' and the 'Action' dropdown is set to 'Add'. There are two input fields: 'Start IP Address' with the value '10.1.0.250' and 'End IP Address' with the value '10.1.0.254' and a '(optional)' label. At the bottom right, there are 'Apply' and 'Revert' buttons.

IP Service > DHCP > Server	
Step:	2. Configure Excluded Address
Action:	Add
Start IP Address	10.1.0.250
End IP Address	10.1.0.254 (optional)
<div>Apply Revert</div>	

To show the IP addresses excluded for DHCP clients:

1. Click IP Service, DHCP, Server.
2. Select Configure Excluded Addresses from the Step list.
3. Select Show from the Action list.

Figure 431: Showing Excluded Addresses on the DHCP Server

IP Service > DHCP > Server

Step: 2. Configure Excluded Address Action: Show

DHCP Excluded Address List Total: 1

	Start IP Address	End IP Address
☐	10.1.0.250	10.1.0.254

Delete Revert

Configuring Address Pools

Use the IP Service > DHCP > Server (Configure Pool – Add) page configure IP address pools for each IP interface that will provide addresses to attached clients via the DHCP server.

Command Usage

- ◆ First configure address pools for the network interfaces. Then you can manually bind an address to a specific client if required. However, note that any static host address must fall within the range of an existing network address pool. You can configure up to 8 network address pools, and up to 32 manually bound host address pools (i.e., one address per host pool). Just note that any address specified in a host address pool must fall within the range of a configured network address pool.
- ◆ When a client request is received, the switch first checks for a network address pool matching the gateway where the request originated (i.e., if the request was forwarded by a relay server). If there is no gateway in the client request (i.e., the request was not forwarded by a relay server), the switch searches for a network pool matching the interface through which the client request was received. It then searches for a manually configured host address that falls within the matching network pool. If no manually configured host address is found, it assigns an address from the matching network address pool. However, if no matching address pool is found the request is ignored.
- ◆ When searching for a manual binding, the switch compares the client identifier and then the hardware address for DHCP clients. Since BOOTP clients cannot transmit a client identifier, you must configure a hardware address for this host type. If no manual binding has been specified for a host entry with a hardware address or client identifier, the switch will assign an address from the first matching network pool.
- ◆ If the subnet mask is not specified for network or host address pools, the class A, B, or C natural mask is used (see [“Specifying Network Interfaces” on page 697](#)). The DHCP server assumes that all host addresses are available. You can exclude subsets of the address space by using the IP Service > DHCP > Server (Configure Excluded Addresses – Add) page.

Parameters

These parameters are displayed:

Creating a New Address Pool

- ◆ **Pool Name** – A string or integer. (Range: 1-32 characters)
- ◆ **Type** – Sets the address pool type to Network or Host.

Setting Parameters for a Network Pool

- ◆ **IP** – The IP address of the DHCP address pool.
- ◆ **Subnet Mask** – The bit combination that identifies the network (or subnet) and the host portion of the DHCP address pool.

Setting Parameters for a Static Host

- ◆ **IP** – The IP address to assign to the host.
- ◆ **Subnet Mask** – Specifies the network mask of the client.
- ◆ **Client-Identifier** – A unique designation for the client device, either a text string (1-15 characters) or hexadecimal value. The information included in the identifier is based on RFC 2132 Option 60, and must be unique for all clients in the same administrative domain.
- ◆ **Hardware Address** – Specifies the MAC address and protocol used on the client. (Options: Ethernet, IEEE802, FDDI, None; Default: Ethernet)

Setting Optional Parameters

- ◆ **Default Router** – The IP address of the primary and alternate gateway router. The IP address of the router should be on the same subnet as the client.
- ◆ **DNS Server** – The IP address of the primary and alternate DNS server. DNS servers must be configured for a DHCP client to map host names to IP addresses.
- ◆ **Netbios Server** – IP address of the primary and alternate NetBIOS Windows Internet Naming Service (WINS) name server used for Microsoft DHCP clients.
- ◆ **Netbios Type** – NetBIOS node type for Microsoft DHCP clients. (Options: Broadcast, Hybrid, Mixed, Peer to Peer; Default: Hybrid)
- ◆ **Domain Name** – The domain name of the client. (Range: 1-128 characters)
- ◆ **Bootfile** – The default boot image for a DHCP client. This file should be placed on the Trivial File Transfer Protocol (TFTP) server specified as the Next Server.
- ◆ **Next Server** – The IP address of the next server in the boot process, which is typically a Trivial File Transfer Protocol (TFTP) server.

- ◆ **Lease Time** – The duration that an IP address is assigned to a DHCP client.
(Options: Finite, Infinite; Default: Infinite)

The lease time configured by this attribute will be overridden by a DHCP client request that includes an Option 51 lease time.

Web Interface

To configure DHCP address pools:

1. Click IP Service, DHCP, Server.
2. Select Configure Pool from the Step list.
3. Select Add from the Action list.
4. Set the pool Type to Network or Host.
5. Enter the IP address and subnet mask for a network pool or host. If configuring a static binding for a host, enter the client identifier or hardware address for the host device. Configure the optional parameters such as a gateway server and DNS server.
6. Click Apply.

Figure 432: Configuring DHCP Server Address Pools (Network)

IP Service > DHCP > Server

Step: 3. Configure Pool Action: Add

Pool Name: tps

Type: Network

IP: 10.1.0.0

Subnet Mask: 255.255.255.0

<<Option>>

Default Router: 10.1.0.253

DNS Server: 10.2.3.4

Netbios Server: 10.1.0.33

Netbios Type: Hybrid

Bootfile: wme.bat

Lease Time: Infinite

Default Router 2:

DNS Server 2:

Netbios Server 2:

Domain Name: example.com

Next Server: 10.1.0.21

Apply Revert

Figure 433: Configuring DHCP Server Address Pools (Host)

IP Service > DHCP > Server

Step: 3. Configure Pool Action: Add

Pool Name: mgr
Type: Host
IP: 10.1.0.19
Subnet Mask: 255.255.255.0
Client Identifier: Text bear
Hardware Address: Ethernet 00-e0-29-94-34-28

<<Option>>

Default Router: 10.1.0.253
DNS Server: 10.2.3.4
Netbios Server: 10.1.0.33
Netbios Type: Hybrid
Bootfile: wme.bat
Lease Time: Infinite

Default Router 2:
DNS Server 2:
Netbios Server 2:
Domain Name: example.com
Next Server: 10.1.0.21

To show the configured DHCP address pools:

1. Click IP Service, DHCP, Server.
2. Select Configure Pool from the Step list.
3. Select Show from the Action list.

Figure 434: Showing Configured DHCP Server Address Pools

IP Service > DHCP > Server

Step: 3. Configure Pool Action: Show

DHCP Pool List Total: 2

	Pool Name	Pool Type	IP	Subnet Mask
☐	mgr	Host	10.1.0.19	255.255.255.0
☐	tps	Network	10.1.0.0	255.255.255.0

Delete Revert

Displaying Address Bindings

Use the IP Service > DHCP > Server (Show IP Binding) page display the host devices which have acquired an IP address from this switch's DHCP server.

Parameters

These parameters are displayed:

- ◆ **IP Address** – IP address assigned to host.


- ◆ **MAC Address** – MAC address of host.
- ◆ **Lease Time** – Duration that this IP address can be used by the host.
- ◆ **Start Time** – Time this address was assigned by the switch.

Web Interface

To show the addresses assigned to DHCP clients:

1. Click IP Service, DHCP, Server.
2. Select Show IP Binding from the Step list.

Figure 435: Shows Addresses Assigned by the DHCP Server



IP Service > DHCP > Server				
Step: 4. Show IP Binding				
IP Binding List Total: 1				
	IP Address	MAC Address	Lease Time	Start Time
☐	192.168.2.61	00-60-6E-00-5F-A1	Infinite	Dec 7 12:20:50 2009

Delete Revert

Forwarding UDP Service Requests

This section describes how this switch can forward UDP broadcast packets originating from host applications to another part of the network when an local application server is not available.

Command Usage

- ◆ Network hosts occasionally use UDP broadcasts to determine information such as address configuration, and domain name mapping. These broadcasts are confined to the local subnet, either as an all hosts broadcast (all ones broadcast - 255.255.255.255), or a directed subnet broadcast (such as 10.10.10.255). To reduce the number of application servers deployed in a multi-segment network, UDP helper can be used to forward broadcast packets for specified UDP application ports to remote servers located in another network segment.
- ◆ To configure UDP helper, enable it globally (see [“Configuring General DNS Service Parameters” on page 637](#)), specify the UDP destination ports for which broadcast traffic will be forwarded (see [“Specifying UDP Destination Ports” on page 654](#)), and specify the remote application servers or the subnet where the servers are located (see [“Specifying the Target Server or Subnet” on page 655](#)).

Enabling the UDP Helper Use the IP Service > UDP Helper > General page to enable the UDP helper globally on the switch.

Parameters

These parameters are displayed:


- ◆ **UDP Helper Status** – Enables or disables the UDP helper. (Default: Disabled)

Web Interface

To enable the UDP help:

1. Click IP Service, UDP Helper, General.
2. Mark the Enabled check box.
3. Click Apply.

Figure 436: Enabling the UDP Helper



Specifying UDP Destination Ports Use the IP Service > UDP Helper > Forwarding page to specify the UDP destination ports for which broadcast traffic will be forwarded when the UDP helper is enabled.

Command Usage

Up to 100 UDP ports can be specified with this command for forwarding to one or more remote servers.

Parameters

These parameters are displayed:

- ◆ **Destination UDP Port** – UDP application port for which UDP service requests are forwarded. (Range: 1-65535)

The following UDP ports are included in the forwarding list when the UDP helper is enabled, and a remote server address is configured:

BOOTP client port 67
BOOTP server port 68
Domain Name Service port 53
IEN-116 Name Service port 42
NetBIOS Datagram Server port 138
NetBIOS Name Server port 137
NTP port 37
TACACS service port 49

TFTP port 69

Web Interface

To specify UDP destination ports for forwarding:

1. Click IP Service, UDP Helper, Forwarding.
2. Select Add from the Action list.
3. Enter a destination UDP port number for which service requests are to be forwarded to a remote application server.
4. Click Apply.

Figure 437: Specifying UDP Destination Ports

The screenshot shows the 'IP Service > UDP Helper > Forwarding' page. At the top, there is a breadcrumb trail. Below it, the 'Action:' dropdown menu is set to 'Add'. Underneath, the 'Destination UDP Port (1-65535)' text input field contains the number '547'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show the configured UDP destination ports:

1. Click IP Service, UDP Helper, Forwarding.
2. Select Show from the Action list.

Figure 438: Showing the UDP Destination Ports

The screenshot shows the 'IP Service > UDP Helper > Forwarding' page with the 'Action:' dropdown menu set to 'Show'. Below this, there is a section titled 'UDP Helper Forwarding Port List' with a 'Total: 1' indicator. A table displays the configured ports:

	Destination UDP Port
☐	547

At the bottom right, there are 'Delete' and 'Revert' buttons.

Specifying the Target Server or Subnet

Use the IP Service > UDP Helper > Address page to specify the application server or subnet (indicated by a directed broadcast address) to which designated UDP broadcast packets are forwarded.

Command Usage

- ◆ Up to 20 helper addresses can be specified.
- ◆ To forward UDP packets with the UDP helper, the clients must be connected to the selected interface, and the interface configured with an IP address.

- ◆ The UDP packets to be forwarded must be specified in the IP Service > UDP Helper > Forwarding page, and the packets meet the following criteria:
 - The MAC address of the received frame must be the all-ones broadcast address (ffff.ffff.ffff).
 - The IP destination address must be one of the following:
 - all-ones broadcast (255.255.255.255)
 - subnet broadcast for the receiving interface
 - The IP time-to-live (TTL) value must be at least 2.
 - The IP protocol must be UDP (17).
 - The UDP destination port must be TFTP, Domain Name System (DNS), Time, NetBIOS, BOOTP or DHCP packet, or a UDP port specified on the IP Service > UDP Helper > Forwarding page.
- ◆ If a helper address is specified on this configuration page, but no UDP ports have been specified on the IP Service > UDP Helper > Forwarding page, broadcast traffic for several UDP protocol types will be forwarded by default as described on [page 654](#).

Parameters

These parameters are displayed:

- ◆ **VLAN ID** – VLAN identifier (Range: 1-4094)
- ◆ **IP Address** – Host address or directed broadcast address to which UDP broadcast packets are forwarded. (Range: 1-65535)

Web Interface

To specify the target server or subnet for forwarding UDP request packets:

1. Click IP Service, UDP Helper, Address.
2. Select Add from the Action list.
3. Enter the address of the remote server or subnet where UDP request packets are to be forwarded.
4. Click Apply.

Figure 439: Specifying the Target Server or Subnet for UDP Requests

IP Service > UDP Helper > Address

Action: Add

VLAN ID: 1

IP Address: 192.168.2.255

Apply Revert

To show the target server or subnet for UDP requests:

1. Click IP Service, UDP Helper, Address.
2. Select Show from the Action list.

Figure 440: Showing the Target Server or Subnet for UDP Requests

IP Service > UDP Helper > Address

Action: Show

VLAN ID: 1

UDP Helper Address List Total: 1

	IP Address
☒	192.168.2.255

Delete Revert

Configuring the PPPoE Intermediate Agent

This section describes how to configure the PPPoE Intermediate Agent (PPPoE IA) relay parameters required for passing authentication messages between a client and broadband remote access servers.

Configuring PPPoE IA Global Settings

Use the IP Service > PPPoE Intermediate Agent (Configure Global) page to enable the PPPoE IA on the switch, set the access node identifier, and set the generic error message.

Command Usage

When PPPoE IA is enabled, the switch inserts a tag identifying itself as a PPPoE IA residing between the attached client requesting network access and the ports connected to broadband remote access servers (BRAS). The switch extracts access-loop information from the client's PPPoE Active Discovery Request, and forwards this information to all trusted ports (designated on the Configure Interface page). The BRAS detects the presence of the subscriber's circuit-ID tag inserted by the switch during the PPPoE discovery phase, and sends this tag as a NAS-port-ID attribute in PPP authentication and AAA accounting requests to a RADIUS server.

Parameters

These parameters are displayed:

- ◆ **PPPoE IA Global Status** – Enables the PPPoE Intermediate Agent globally on the switch. (Default: Disabled)

Note that PPPoE IA must be enabled globally before it can be enabled on an interface.

- ◆ **Access Node Identifier** – String identifying this switch as an PPPoE IA to the PPPoE server. (Range: 1-48 ASCII characters; Default: IP address of first IPv4 interface on the switch.)

The switch uses the access-node-identifier to generate the circuit-id for PPPoE discovery stage packets sent to the BRAS, but does not modify the source or destination MAC address of these PPPoE discovery packets. These messages are forwarded to all trusted ports designated on the Configure Interface page.

- ◆ **Operational Access Node Identifier** – The configured access node identifier.
- ◆ **Generic Error Message** – An error message notifying the sender that the PPPoE Discovery packet was too large. (Range: 0-127; Default: PPPoE Discover packet too large to process. Try reducing the number of tags added.)
- ◆ **Operational Generic Error Message** – The configured generic error message.

Web Interface

To configure global settings for PPPoE IA:

1. Click IP Service, PPPoE Intermediate Agent.
2. Select Configure Global from the Step list.
3. Enable the PPPoE IA on the switch, set the access node identifier, and set the generic error message.
4. Click Apply.

Figure 441: Configuring Global Settings for PPPoE Intermediate Agent

IP Service > PPPoE Intermediate Agent

Step: 1. Configure Global

PPPoE IA Global Status	☐ Enabled
Access Node Identifier	<input type="text"/>
Operational Access Node Identifier	10.254.254.1
Generic Error Message	<input type="text"/>
Operational Generic Error Message	PPPoE Discover packet too large to process. Try reducing the number of tags added.

Apply Revert

Configuring PPPoE IA Interface Settings Use the IP Service > PPPoE Intermediate Agent (Configure Interface) page to enable PPPoE IA on an interface, set trust status, enable vendor tag stripping, and set the circuit ID and remote ID.

Parameters

These parameters are displayed:

- ◆ **Interface** – Port or trunk selection.
- ◆ **PPPoE IA Status** – Enables the PPPoE IA on an interface. (Default: Disabled)
Note that PPPoE IA must also be enabled globally on the switch for this command to take effect.
- ◆ **Trust Status** – Sets an interface to trusted mode to indicate that it is connected to a PPPoE server. (Default: Disabled)
 - Set any interfaces connecting the switch to a PPPoE Server as trusted. Interfaces that connect the switch to users (PPPoE clients) should be set as untrusted.
 - At least one trusted interface must be configured on the switch for the PPPoE IA to function.
- ◆ **Vendor Tag Strip** – Enables the stripping of vendor tags from PPPoE Discovery packets sent from a PPPoE server. (Default: Disabled)
This parameter only applies to trusted interfaces. It is used to strip off vendor-specific tags (which carry subscriber and line identification information) in PPPoE Discovery packets received from an upstream PPPoE server before forwarding them to a user.
- ◆ **Circuit ID** – String identifying the circuit identifier (or interface) on this switch to which the user is connected. (Range: 1-10 ASCII characters; Default: Unit/Port:VLAN-ID, or 0/Trunk-ID:VLAN-ID)
 - The PPPoE server extracts the Line-ID tag from PPPoE discovery stage messages, and uses the Circuit-ID field of that tag as a NAS-Port-ID attribute in AAA access and accounting requests.
 - The switch intercepts PPPoE discovery frames from the client and inserts a unique line identifier using the PPPoE Vendor-Specific tag (0x0105) to PPPoE Active Discovery Initiation (PADI) and Request (PADR) packets. The switch then forwards these packets to the PPPoE server. The tag contains the Line-ID of the customer line over which the discovery packet was received, entering the switch (or access node) where the intermediate agent resides.
 - Outgoing PAD Offer (PADO) and Session-confirmation (PADS) packets sent from the PPPoE Server include the Circuit-ID tag inserted by the switch, and should be stripped out of PADO and PADS packets which are to be passed directly to end-node clients.
- ◆ **Operational Circuit ID** – The configured circuit identifier.

◆ **Remote ID** – String identifying the remote identifier (or interface) on this switch to which the user is connected. (Range: 1-63 ASCII characters; Default: Port MAC address)

◆ **Operational Remote ID** – The configured circuit identifier.

Web Interface

To configure interface settings for PPPoE IA:

1. Click IP Service, PPPoE Intermediate Agent.
2. Select Configure Interface from the Step list.
3. Select Port or Trunk interface type.
4. Enable PPPoE IA on an interface, set trust status, enable vendor tag stripping if required, and set the circuit ID and remote ID.
5. Click Apply.

Figure 442: Configuring Interface Settings for PPPoE Intermediate Agent

Port	PPPoE IA Status	Trust Status	Vendor Tag Strip	Circuit ID	Operation Circuit ID	Remote ID	Operation Remote ID
1	☒ Enabled	☐ Enabled	☐ Enabled		1/1:vid		00-E0-0C-00-00-FE
2	☒ Enabled	☐ Enabled	☐ Enabled		1/2:vid		00-E0-0C-00-00-FF
3	☒ Enabled	☐ Enabled	☐ Enabled		1/3:vid		00-E0-0C-00-01-00
4	☒ Enabled	☐ Enabled	☐ Enabled		1/4:vid		00-E0-0C-00-01-01
5	☒ Enabled	☐ Enabled	☐ Enabled		1/5:vid		00-E0-0C-00-01-02

Showing PPPoE IA Statistics Use the IP Service > PPPoE Intermediate Agent (Show Statistics) page to show statistics on PPPoE IA protocol messages.

Parameters

These parameters are displayed:

- ◆ **Interface** – Port or trunk selection.
- ◆ **Received** – Received PPPoE active discovery messages.
 - **All** – All PPPoE active discovery message types.
 - **PADI** – PPPoE Active Discovery Initiation messages.

- **PADO** – PPPoE Active Discovery Offer messages.
- **PADR** – PPPoE Active Discovery Request messages.
- **PADS** – PPPoE Active Discovery Session-Confirmation messages.
- **PADT** – PPPoE Active Discovery Terminate messages.
- ◆ **Dropped** – Dropped PPPoE active discovery messages.
 - **Response from untrusted** – Response from an interface which not been configured as trusted.
 - **Request towards untrusted** – Request sent to an interface which not been configured as trusted.
 - **Malformed** – Corrupted PPPoE message.

Web Interface

To show statistics for PPPoE IA protocol messages:

1. Click IP Service, PPPoE Intermediate Agent.
2. Select Show Statistics from the Step list.
3. Select Port or Trunk interface type.

Figure 443: Showing PPPoE Intermediate Agent Statistics

IP Service > PPPoE Intermediate Agent

Step: 3. Show Statistics

Interface ☒ Port ☐ Trunk

PPPoE Intermediate Agent Statistics Total: 12

	Port	Received						Dropped		
		All	PADI	PADO	PADR	PADS	PADT	Response from untrusted	Request towards untrusted	Malformed
☐	1	0	0	0	0	0	0	0	0	0
☐	2	0	0	0	0	0	0	0	0	0
☐	3	0	0	0	0	0	0	0	0	0
☐	4	0	0	0	0	0	0	0	0	0
☐	5	0	0	0	0	0	0	0	0	0

General IP Routing

This chapter provides information on network functions including:

- ◆ [Ping](#) – Sends ping message to another node on the network.
- ◆ [Trace Route](#) – 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 162](#)), assign all ports that belong to the same group to these VLANs ([page 165](#)), and then assign an IP interface to each VLAN ([page 613](#) or [page 618](#)). 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 444: 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 613](#). 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 618](#)).

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 675](#)), or use RIP, OSPF, OSPFv3¹⁶, or BGPv4¹⁶ dynamic routing protocols 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 675](#), 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 619](#).

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 637](#)), and one or more DNS servers specified (see [“Configuring a List of Name Servers” on page 640](#), or [“Configuring Static DNS Host to Address Entries” on page 641](#)).
- ◆ **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 445: Pinging a Network Device

The screenshot shows the 'IP > General > Ping' configuration page. It includes three input fields: 'Host Name/IP Address' (empty), 'Probe Count (1-16)' (set to 5), and 'Packet Size (IPv4 : 32-1472, IPv6 : 0-512)' (empty). Below these fields are 'Apply' and 'Revert' buttons. A 'Result' section displays the output of a ping command: 'PING to 192.168.0.3, by 5 of 32-byte payload ICMP packets, timeout is 3 seconds'. The output shows five response times: 10 ms, 0 ms, 0 ms, 0 ms, and 0 ms. Below this, it provides statistics for 192.168.0.3: '5 packets transmitted, 5 packets received (100%), 0 packets lost (0%)'. It also shows approximate round trip times: 'Minimum = 0 ms, Maximum = 10 ms, Average = 2 ms'.

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.

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 446: Tracing the Route to a Network Device



Address Resolution Protocol

If IP routing is enabled (page 691), 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.

Basic ARP Configuration

Use the IP > ARP (Configure General) page to specify the timeout for ARP cache entries, or to enable Proxy ARP for specific VLAN interfaces.


Command Usage

Proxy ARP

When a node in the attached subnetwork does not have routing or a default gateway configured, Proxy ARP can be used to forward ARP requests to a remote subnetwork. When the router receives an ARP request for a remote network and Proxy ARP is enabled, it determines if it has the best route to the remote network, and then answers the ARP request by sending its own MAC address to the

requesting node. That node then sends traffic to the router, which in turn uses its own routing table to forward the traffic to the remote destination.

Figure 447: Proxy ARP



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.

- ◆ **Proxy ARP** – Enables or disables Proxy ARP for specified VLAN interfaces, allowing a non-routing device to determine the MAC address of a host on another subnet or network. (Default: Disabled)

End stations that require Proxy ARP must view the entire network as a single network. These nodes must therefore use a smaller subnet mask than that used by the router or other relevant network devices.

Extensive use of Proxy ARP can degrade router performance because it may lead to increased ARP traffic and increased search time for larger ARP address tables.

Web Interface

To configure the timeout for the ARP cache or to enable Proxy ARP for a VLAN (i.e., IP subnetwork):

1. Click IP, ARP.
2. Select Configure General from the Step List.
3. Set the timeout to a suitable value for the ARP cache, or enable Proxy ARP for subnetworks that do not have routing or a default gateway.
4. Click Apply.

Figure 448: Configuring General Settings for ARP

The screenshot shows a web-based configuration interface for ARP settings. At the top, it says "IP > ARP". Below that, there's a "Step:" dropdown menu set to "1. Configure General". The main configuration area includes a "Timeout (300-86400)" field with the value "1200" and a "sec" unit. Below this is a "Proxy ARP" section with a "VLAN" dropdown menu set to "1" and a "Status" checkbox labeled "Enabled" which is currently unchecked. At the bottom right, there are "Apply" and "Revert" buttons.

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 an 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 449: Configuring Static ARP Entries

The screenshot shows the 'IP > ARP' configuration page. The 'Step' dropdown is set to '2. Configure Static Address' and the 'Action' dropdown is set to 'Add'. Below these, 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 450: Displaying Static ARP Entries

The screenshot shows the 'IP > ARP' configuration page with the 'Action' dropdown set to 'Show'. Below the configuration fields, there is a table titled 'Static Address List' with a 'Total: 2' indicator. The table has four columns: a checkbox, 'IP Address', 'MAC Address', and 'Interface'. It contains two entries. At the bottom right, there are 'Delete' and 'Revert' buttons.

	IP Address	MAC Address	Interface
☐	10.2.78.105	00-E0-0C-00-00-FD	VLAN 1
☐	10.2.78.254	00-12-D9-75-DC-5B	VLAN 1

Displaying Dynamic or Local ARP Entries


Use the IP > ARP (Show Information – ARP Addresses) page to display dynamic or local entries in the ARP cache. 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.

Web Interface

To display all dynamic and local entries in the ARP cache:

1. Click IP, ARP.
2. Select Show Information from the Step List.
3. Click ARP Addresses.

Figure 451: Displaying ARP Entries



IP Address	MAC Address	Type	Interface
192.168.0.3	70-72-CF-83-34-66	other	VLAN 1
192.168.0.99	00-60-6E-00-5F-A1	dynamic	VLAN 1

Displaying ARP Statistics

Use the IP > ARP (Show Information – ARP Statistics) 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.
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 452: Displaying ARP Statistics

The screenshot shows the 'IP > ARP' configuration page. At the top, there is a 'Step:' dropdown menu set to '3. Show Information'. Below this, there are two radio buttons: 'ARP Address' (unselected) and 'Statistics' (selected). Under the 'Statistics' section, the title 'ARP Statistics' is displayed. Below the title, there are four statistics listed in two columns: 'Received Request' with a value of 2, 'Received Reply' with a value of 1, 'Sent Request' with a value of 1, and 'Sent Reply' with a value of 2. At the bottom right of the statistics section, there is a 'Refresh' button.

ARP Statistics			
Received Request	2	Sent Request	1
Received Reply	1	Sent Reply	2

Configuring Static Routes

This router can dynamically configure routes to other network segments using dynamic routing protocols (i.e., RIP, OSPF, BGP). 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 512 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 678).
- ◆ 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 (see page 701 or 728 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, 120 for RIP, 20 for eBGP, and 200 for iBGP. (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 453: Configuring Static Routes

IP > Routing > Static Routes

Action: Add ▼

Destination IP Address	10.2.48.0
Net Mask / Prefix Length	255.255.255.0
Next Hop	10.2.48.1
Distance (1-255)	5 (Optional)

Apply Revert

To display static routes:

1. Click IP, Routing, Static Routes.
2. Select Show from the Action List.

Figure 454: Displaying Static Routes

IP > Routing > Static Routes

Action:

Static Table List Total: 3

☐	Destination IP Address	Net Mask / Prefix Length	Next Hop	Distance
☐	10.2.48.2	255.255.255.0	10.2.48.1	5
☐	10.5.36.123	255.255.0.0	10.5.36.1	2
☐	5566::	60	7788::12	1

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 (and the “show ip route” command described in the *CLI Reference Guide*) 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., configured 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 455: Displaying the Routing Table

IP > Routing > Routing Table					
Action: Show Information					
Routing Table List Total: 5					
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 675](#)), or dynamically generated by OSPFv2 (see [“Configuring the Open Shortest Path First Protocol \(Version 2\)” on page 710](#)), 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 456: Setting the Maximum ECMP Number



The screenshot shows a web interface for configuring the maximum ECMP number. The breadcrumb navigation at the top reads "IP > Routing > Routing Table". Below this, there is a section titled "Action:" with a dropdown menu set to "Configure ECMP Number". Underneath, the label "ECMP Number (1-8)" is followed by a text input field containing the value "4". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

Configuring Router Redundancy

Router redundancy protocols use a virtual IP address to support a primary router and multiple backup routers. The backup routers can be configured to take over the workload if the master router fails, or can also be configured to share the traffic load. The primary goal of router redundancy 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.

This switch supports the Virtual Router Redundancy Protocol (VRRP). VRRP allows you to specify the interface of one of the routers participating in the virtual group as the address for the master virtual router, or to configure an arbitrary address for the virtual master router. VRRP then selects the backup routers based on the specified virtual router priority.

Router redundancy can be set up in any of the following configurations. These examples use the address of one of the participating routers as the master router. When the virtual router IP address is not a real address, the master router is selected based on priority. When the priority is the same on several competing routers, then the router with the highest IP address is selected as the master.

Figure 457: Master Virtual Router with Backup Routers




Figure 458: Several Virtual Master Routers Using Backup Routers





Figure 459: Several Virtual Master Routers Configured for Mutual Backup and Load Sharing



Note: Load sharing can be accomplished by assigning a subset of addresses to different host address pools using the DHCP server. (See [“Configuring Address Pools”](#) on page 649)

Configuring VRRP Groups

Use the IP > VRRP pages to configure VRRP. To configure VRRP groups, select an interface on each router in the group that will participate in the protocol as the master router or a backup router. To select a specific device as the master router, set the address of this interface as the virtual router address for the group. Now set the same virtual address and a priority on the backup routers, and configure an authentication string. You can also enable the preempt feature which allows a router to take over as the master router when it comes on line if it has a higher priority than the currently active master router.

Command Usage

Address Assignment –

- ◆ To designate a specific router as the VRRP master, the IP address assigned to the virtual router must already be configured on the router that will become the Owner of the group address. In other words, the IP address for the virtual router exists on one, and only one, router in the virtual router group, and the network mask for the virtual router address is derived from the Owner. The Owner will also assume the role of the Master virtual router in the group.
- ◆ If a virtual address is assigned to the group which does not exist on any of the group members, then the master router is selected based on priority. In cases where the configured priority is the same on several group members, then the master router with the highest IP address is selected from this group.
- ◆ If you have multiple secondary addresses configured on the current VLAN interface, you can add any of these addresses to the virtual router group.
- ◆ The interfaces of all routers participating in a virtual router group must be within the same IP subnet.
- ◆ VRRP creates a virtual MAC address for the master router based on a standard prefix, with the last octet equal to the group ID. When a backup router takes over as the master, it continues to forward traffic addressed to this virtual MAC address. However, the backup router cannot reply to ICMP pings sent to addresses associated with the virtual group because the IP address owner is off line.

Virtual Router Priority –

- ◆ The Owner of the virtual IP address is automatically assigned the highest possible virtual router priority of 255. The backup router with the highest priority will become the master router if the current master fails. However, because the priority of the virtual IP address Owner is the highest, the original master router will always become the active master router when it recovers.
- ◆ If two or more routers are configured with the same VRRP priority, the router with the higher IP address is elected as the new master router if the current master fails.

Preempting the Acting Master –

- ◆ The virtual IP Owner has the highest priority, so no other router can preempt it, and it will always resume control as the master virtual router when it comes back on line. The preempt function only allows a backup router to take over from a master router if no router in the group is the virtual IP owner, or from another backup router that is temporarily acting as the group master. If preemption is enabled and this router has a higher priority than the current acting master when it comes on line, it will take over as the acting group master.

- ◆ You can add a delay to the preempt function to give additional time to receive an advertisement message from the current master before taking control. If the router attempting to become the master has just come on line, this delay also gives it time to gather information for its routing table before actually preempting the currently active master router.

Parameters

These parameters are displayed:

Adding a VRRP Group

- ◆ **VRID** – VRRP group identifier. (Range: 1-255)
The maximum number of groups which can be defined is 64.
- ◆ **VLAN** – ID of a VLAN configured with an IP interface. (Range: 1-4094; Default: 1)

Adding a Virtual IP Address

- ◆ **VLAN ID** – ID of a VLAN configured with an IP interface. (Range: 1-4094)
- ◆ **VRID** – VRRP group identifier. (Range: 1-255)
- ◆ **IP Address** – Virtual IP address for this group.
Use the IP address of a real interface on this router to make it the master virtual router for the group. Otherwise, use the virtual address for an existing group to make it a backup router, or to compete as the master based on configured priority if no other members are set as the owner of the group address.

Configuring Detailed Settings

- ◆ **VLAN ID** – VLAN configured with an IP interface. (Range: 1-4094)
- ◆ **VRID** – VRRP group identifier. (Range: 1-255)
- ◆ **Advertisement Interval** – Interval at which the master virtual router sends advertisements communicating its state as the master. (Range: 1-255 seconds; Default: 1 second)

VRRP advertisements from the current master virtual router include information about its priority and current state as the master.

VRRP advertisements are sent to the multicast address 224.0.0.8. Using a multicast address reduces the amount of traffic that has to be processed by network devices that are not part of the designated VRRP group.

If the master router stops sending advertisements, backup routers will bid to become the master router based on priority. The dead interval before attempting to take over as the master is three times the hello interval plus half a second.

- ◆ **Priority** – The priority of this router in a VRRP group. (Range: 1-254; Default: 100)
 - The priority for the VRRP group address owner is automatically set to 255.
 - The priority for backup routers is used to determine which router will take over as the acting master router if the current master fails.
- ◆ **Preempt Mode** – Allows a backup router to take over as the master virtual router if it has a higher priority than the acting master virtual router (i.e., a master router that is not the group's address owner, or another backup router that has taken over from the previous master). (Default: Enabled)
- ◆ **Preempt Delay Time** – Time to wait before issuing a claim to become the master. (Range: 0-120 seconds; 0 seconds)
- ◆ **Authentication Mode** – Authentication mode used to verify VRRP packets received from other routers. (Options: None, Simple Text; Default: None)

If simple text authentication is selected, then you must also enter an authentication string.

All routers in the same VRRP group must be set to the same authentication mode, and be configured with the same authentication string.

Plain text authentication does not provide any real security. It is supported only to prevent a misconfigured router from participating in VRRP.
- ◆ **Authentication String** – Key used to authenticate VRRP packets received from other routers. (Range: 1-8 alphanumeric characters)

When a VRRP packet is received from another router in the group, its authentication string is compared to the string configured on this router. If the strings match, the message is accepted. Otherwise, the packet is discarded.
- ◆ **State** – VRRP router role. (Values: Master, Backup)
- ◆ **Virtual MAC Address** – Virtual MAC address for this group.
- ◆ **Master Router** – The primary router servicing this group.
- ◆ **Master Priority** – The priority of the master router.
- ◆ **Master Advertisement Interval** – The interval at which the master router sends messages advertising itself as the group master.
- ◆ **Master Down Interval** – If no advertisement message is received from the master router after this interval, backup routers will assume that the master is dead, and will start bidding to become the group master.

Web Interface

To configure VRRP:

1. Click IP, VRRP.
2. Select Configure Group ID from the Step List.
3. Select Add from the Action List.
4. Enter the VRID group number, and select the VLAN (i.e., IP subnet) which is to be serviced by this group.
5. Click Apply.

Figure 460: Configuring the VRRP Group ID

The screenshot shows the 'IP > VRRP' configuration page. At the top, there is a 'Step:' dropdown menu set to '1. Configure Group ID' and an 'Action:' dropdown menu set to 'Add'. Below these, there are two input fields: 'VRID (1-255)' with the value '1' and 'VLAN' with a dropdown menu set to '1'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

To show the configured VRRP groups:

1. Click IP, VRRP.
2. Select Configure Group ID from the Step List.
3. Select Show from the Action List.

Figure 461: Showing Configured VRRP Groups

The screenshot shows the 'IP > VRRP' configuration page with the 'Action:' dropdown menu set to 'Show'. Below the configuration fields, there is a table titled 'VRRP Group ID List' with a 'Total: 2' indicator. The table has three columns: a checkbox, 'VRID', and 'VLAN'. There are two rows of data.

	VRID	VLAN
☐	1	1
☐	2	1

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

To configure the virtual router address for a VRRP group:

1. Click IP, VRRP.
2. Select Configure Group ID from the Step List.
3. Select Add IP Address from the Action List.

4. Select a VLAN, a VRRP group identifier, and enter the IP address for the virtual router.
5. Click Apply.

Figure 462: Setting the Virtual Router Address for a VRRP Group

IP > VRRP

Step: 1. Configure Group ID Action: Add IP Address

VLAN ID: 1

VRID: 1

IP Address: 192.168.2.9

Apply Revert

To show the virtual IP address assigned to a VRRP group:

1. Click IP, VRRP.
2. Select Configure Group ID from the Step List.
3. Select Show IP Addresses from the Action List.
4. Select a VLAN, and a VRRP group identifier.

Figure 463: Showing the Virtual Addresses Assigned to VRRP Groups

IP > VRRP

Step: 1. Configure Group ID Action: Show IP Addresses

VLAN ID: 1

VRID: 1

VRRP Group IP List Total: 1

	IP Address
☐	192.168.2.9

Delete Revert

To configure detailed settings for a VRRP group:

1. Click IP, VRRP.
2. Select Configure Group ID from the Step List.
3. Select Configure Detail from the Action List.
4. Select a VRRP group identifier, and set any of the VRRP protocol parameters as required.
5. Click Apply.

Figure 464: Configuring Detailed Settings for a VRRP Group

The screenshot shows the 'IP > VRRP' configuration page. At the top, there are two dropdown menus: 'Step: 1. Configure Group ID' and 'Action: Configure Detail'. Below these, the configuration parameters are listed in a table-like format:

VLAN ID	1
VRID	1
Advertisement Interval (1-255)	1 sec
Priority (1-254)	255
Preempt Mode	☒ Enabled
Preempt Delay Time (0-120)	0 sec
Authentication Mode	None
Authentication String	
State	Master
Virtual MAC Address	00-00-5E-00-01-01
Master Router	192.168.0.2
Master Priority	255
Master Advertisement Interval	1 sec
Master Down Interval	3

Displaying VRRP Global Statistics

Use the IP > VRRP (Show Statistics – Global Statistics) page to display counters for errors found in VRRP protocol packets.

Parameters

These parameters are displayed:

- ◆ **VRRP Packets with Invalid Checksum** – The total number of VRRP packets received with an invalid VRRP checksum value.
- ◆ **VRRP Packets with Unknown Error** – The total number of VRRP packets received with an unknown or unsupported version number.
- ◆ **VRRP Packets with Invalid VRID** – The total number of VRRP packets received with an invalid VRID for this virtual router.

Web Interface

To show counters for errors found in VRRP protocol packets:

1. Click IP, VRRP.
2. Select Show Statistics from the Step List.
3. Click Global Statistics.

Figure 465: Showing Counters for Errors Found in VRRP Packets

IP > VRRP

Step: 2. Show Statistics

☒ Global Statistics
☐ Group Statistics

Global Statistics

VRRP Packets with Invalid Checksum 0

VRRP Packets with Unknown Error 3

VRRP Packets with Invalid VRID 0

Displaying VRRP Group Statistics

Use the IP > VRRP (Show Statistics – Group Statistics) page to display counters for VRRP protocol events and errors that have occurred on a specific VRRP interface.

Parameters

These parameters are displayed:

- ◆ **VLAN ID** – VLAN configured with an IP interface. (Range: 1-4094)
- ◆ **VRID** – VRRP group identifier. (Range: 1-255)

The following statistics are displayed:

Table 43: VRRP Group Statistics

Parameter	Description
Times Transitioned to Master	Number of times this router has transitioned to master.
Received Advertisement Packets	Number of VRRP advertisements received by this router.
Received Error Advertisement Interval Packets	Number of VRRP advertisements received for which the advertisement interval is different from the one configured for the local virtual router.
Received Authentication Failure Packets	Number of VRRP packets received that do not pass the authentication check.
Received Error IP TTL VRRP Packets	Number of VRRP packets received by the virtual router with IP TTL (Time-To-Live) not equal to 255.
Received Priority 0 VRRP Packets	Number of VRRP packets received by the virtual router with priority set to 0.
Sent Priority 0 VRRP Packets	Number of VRRP packets sent by the virtual router with priority set to 0. A priority value of zero indicates that the group master has stopped participating in VRRP, and is used to quickly transition a backup unit to master mode without having to wait for the master to time out.
Received Invalid Type VRRP Packets	Number of VRRP packets received by the virtual router with an invalid value in the “type” field.

Table 43: VRRP Group Statistics (Continued)

Parameter	Description
Received Error Address List VRRP Packets	Number of packets received for which the address list does not match the locally configured list for the virtual router.
Received Invalid Authentication Type VRRP Packets	Number of packets received with an unknown authentication type.
Received Mismatch Authentication Type VRRP Packets	Number of packets received with "Auth Type" not equal to the locally configured authentication method.
Received Error Packets Length VRRP Packets	Number of packets received with a packet length less than the length of the VRRP header.

Web Interface

To show counters for VRRP protocol events and errors that occurred on a specific VRRP interface:

1. Click IP, VRRP.
2. Select Show Statistics from the Step List.
3. Click Group Statistics.

Figure 466: Showing Counters for Errors Found in a VRRP Group

IP > VRRP

Step: 2. Show Statistics

Group Statistics

VLAN ID	1
VRID	1
Times Transitioned to Master	1
Received Advertisement Packets	0
Received Error Advertisement Interval Packets	0
Received Authentication Failure Packets	0
Received Error IP TTL VRRP Packets	0
Received Priority 0 VRRP Packets	0
Sent Priority 0 VRRP Packets	0
Received Invalid Type VRRP Packets	0
Received Error Address List VRRP Packets	0
Received Invalid Authentication Type VRRP Packets	0
Received Mismatch Authentication Type VRRP Packets	0
Received Error Packets Length VRRP Packets	0

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 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 467: 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 704](#)).

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 704](#)).

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 697](#). (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 704](#)) 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 704](#)) 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 701](#)). 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: 1-736; Default: 736)
- ◆ **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 703](#)) 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 468: Configuring General Settings for RIP

The screenshot shows the 'Routing Protocol > RIP > General' configuration page. At the top, there is a breadcrumb trail and an 'Action:' dropdown menu set to 'Configure'. The page is divided into two main sections: 'Global' and 'Basic Timer'. The 'Global' section contains the following settings: 'RIP Routing Process' (checked 'Enabled'), 'Global RIP Version' (set to 'By interface'), 'RIP Default Metric (1-15)' (set to '1'), 'RIP Max Prefix (1-11766)' (set to '11766'), 'Default Information Originate' (unchecked 'Enabled'), 'Default Distance (1-255)' (set to '120'), 'Number of Route Changes' (set to '0'), and 'Number of Queries' (set to '0'). The 'Basic Timer' section contains: 'Update (5-2147483647)' (set to '30' sec), 'Timeout (90-360)' (set to '180' sec), and 'Garbage Collection (60-240)' (set to '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 701](#)) 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 469: Clearing Entries from the Routing Table

Routing Protocol > RIP > General

Action: Clear Route

Clear Route by Type ☒ Network

Network IP Address

Prefix Length (1-32)

Apply 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: 1-4094)

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 470: 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 471: Showing Network Interfaces Using RIP

Routing Protocol > RIP > Network

Action: Show

By: ☒ IP Address ☐ VLAN

RIP Network Address List Total: 1

	Subnet Address	Prefix Length
☐	10.1.0.0	16

Delete Revert

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 472: Specifying a Passive RIP Interface

Routing Protocol > RIP > Passive Interface

Action: Add

VLAN ID (1-4094) 1

Apply Revert

To show the passive RIP interfaces:

1. Click Routing Protocol, RIP, Passive Interface.
2. Select Show from the Action list.

Figure 473: Showing Passive RIP Interfaces

Routing Protocol > RIP > Passive Interface

Action: Show

VLAN ID List Total: 1

	VLAN ID
☐	1

Delete Revert

Specifying Static Neighbors

Use the Routing Protocol > RIP > Passive Interface (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 474: 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 475: Showing Static RIP Neighbors

Routing Protocol > RIP > Neighbor Address

Action: Show

Neighbor Address List Total: 1

	IP Address
☐	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 693.)

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 476: Redistributing External Routes into RIP

Routing Protocol > RIP > Redistribute

Action: Add ▼

Protocol: OSPF ▼

Metric (1-16): 3 (Optional)

Apply Revert

To show external routes imported into RIP:

1. Click Routing Protocol, RIP, Redistribute.
2. Select Show from the Action list.

Figure 477: Showing External Routes Redistributed into RIP

Routing Protocol > RIP > Redistribute

Action: Show

Redistribute List Total: 1

	Protocol	Metric
☐	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 478: Setting the Distance Assigned to External Routes

Routing Protocol > RIP > Distance

Action: **Add**

Distance (1-255): 120

IP Address: 192.158.3.0

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 479: Showing the Distance Assigned to External Routes

Routing Protocol > RIP > Distance

Action: **Show**

RIP Distance List Total: 1

	Distance	IP Address	Subnet Mask
☐	120	192.158.3.0	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 693](#)).
- ◆ 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 693.](#))

◆ **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 693.](#))

◆ **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 480: Configuring a Network Interface for RIP

Routing Protocol > RIP > Interface

Action: Add

VLAN ID (1-4094): 1

Send Version: RIPv1 Compatible

Receive Version: RIPv1 and RIPv2

Authentication Type: Simple Password

Authentication Key: mighty

Instability Prevention: Poison Reverse

Apply Revert

To show the network interface settings configured for RIP:

1. Click Routing Protocol, RIP, Interface.
2. Select Show from the Action list.

Figure 481: Showing RIP Network Interface Settings

Routing Protocol > RIP > Interface						
Action: Show						
Interface Settings List Total: 2						
☐	VLAN ID	Send Version	Receive Version	Authentication Type	Authentication Key	Instability Prevention
☐	1	RIPv1 Compatible	RIPv1 and RIPv2	Simple Password	mighty	Poison Reverse
☐	2	RIPv1 Compatible	RIPv1 and RIPv2	No Authentication		Poison Reverse

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 482: 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 483: Showing RIP Peer Information

Routing Protocol > RIP > Statistics				
Action: Show Peer Information				
Peer Information Total: 2				
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 484: 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 485: 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 726](#)). To further reduce the amount of routes passed between areas, an area can be configured as a stub ([page 719](#), [page 723](#)) or a not-so-stubby area ([page 719](#), [page 720](#)).

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 486: 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 726](#).

- ◆ 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 487: Defining OSPF Network Areas Based on Addresses

The screenshot shows the 'Routing Protocol > OSPF > Network Area' configuration page. At the top, there is a breadcrumb trail. Below it, an 'Action:' dropdown menu is set to 'Add'. The main configuration area contains four input fields: '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, there are two buttons: 'Apply' and 'Revert'.

Routing Protocol > OSPF > Network Area	
Action:	Add
Process ID (1-65535)	1
IP Address	192.168.0.0
Netmask	255.255.255.0
Area ID	0.0.0.0
<div>Apply Revert</div>	

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 488: Showing OSPF Network Areas

Routing Protocol > OSPF > Network Area

Action: Show

Process ID: 1

Network Area Address List Total: 2

☐	IP Address	Netmask	Area ID
☐	192.168.0.0	255.255.255.0	0.0.0.0
☐	192.168.1.0	255.255.255.0	192.168.1.0

Delete Revert

To to show the OSPF process identifiers:

1. Click Routing Protocol, OSPF, Network Area.
2. Select Show Process from the Action list.

Figure 489: Showing OSPF Process Identifiers

Routing Protocol > OSPF > Network Area

Action: Show Process

Process List Total: 1

☐	Process ID
☐	1

Delete Revert

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 728](#)). 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 490: 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 675.](#))
 - **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 728.](#))
- ◆ **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 728](#)).

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 491: 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 492: Showing General Settings for OSPF

Routing Protocol > OSPF > System

Action: Show

Process ID: 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 493: Adding an NSSA or Stub

The screenshot shows a web interface for configuring OSPF areas. The breadcrumb trail at the top 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 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 494: 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 495: 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 711](#)). Then create an NSSA as described under [“Adding an NSSA or Stub” on page 719](#).
- ◆ 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 711](#)).
- ◆ **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 728](#)), 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 496: 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	☒ Enabled	☒ 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 497: 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 711](#)). Then create a stub as described under [“Adding an NSSA or Stub” on page 719](#).
- ◆ 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 711](#)).

- ◆ **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 498: Configuring Protocol Settings for a Stub

Routing Protocol > OSPF > Area

Step: 1. Configure Area Action: Configure Stub Area

Process ID: 1

Stub Area List Total: 1

Area ID	Default Cost (0-16777215)	Summary
192.168.3.0	1	Summary

Apply Revert

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 711](#)).
- ◆ **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 499: 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 500: 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 726).
- ◆ 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 711](#)).
- ◆ **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 501: Configuring Route Summaries for an Area Range

The screenshot shows a web interface for configuring OSPF Area Range. The breadcrumb trail at the top is "Routing Protocol > OSPF > Area Range". Below this, there is a section for "Action:" with a dropdown menu set to "Add". The main configuration area contains several fields: "Process ID" with a dropdown set to "1", "Area ID" with a text input containing "192.168.0.0", "Range Network" with a text input containing "192.168.0.0", "Range Netmask" with a text input containing "255.255.0.0", and "Advertising" with a dropdown menu set to "Advertise". At the bottom right of the form, there are two buttons: "Apply" and "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 502: Showing Configured Route Summaries

Routing Protocol > OSPF > Area Range

Action: Show

Process ID: 1

Area Range List Total: 1


☐	Area ID	Range Network	Range Netmask	Advertising
☐	192.168.0.0	192.168.0.0	255.255.0.0	Advertise

Delete Revert

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 503: 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 714).

Parameters

These parameters are displayed:

- ◆ **Process ID** – Process ID as configured in the Network Area configuration screen (see [page 711](#)).

- ◆ **Protocol Type** – Specifies the external routing protocol type for which routing information is to be redistributed into the local routing domain. (Options: BGP, RIP, Static; 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 714](#)).

- ◆ **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 504: Importing External Routes

Routing Protocol > OSPF > Redistribute

Action: Add

Process ID 1

Protocol Type RIP

Metric Type 1

Metric (0-16777214) 3

Tag (0-4294967295) 1

Apply Revert

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 505: Showing Imported External Route Types

Routing Protocol > OSPF > Redistribute

Action: Show

Process ID 1

Redistribute List Total: 1

	Protocol Type	Metric Type	Metric	Tag
☐	RIP	1	3	1

Delete Revert

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 726](#)).
- ◆ 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 711](#)).
- ◆ **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 506: Summarizing External Routes

The screenshot shows a web interface for configuring OSPF summary addresses. The breadcrumb trail at the top is "Routing Protocol > OSPF > Summary Address". Below this, there is a form with the following fields:

- Action:** A dropdown menu currently showing "Add".
- Process ID:** A dropdown menu currently showing "1".
- IP Address:** A text input field containing "192.168.0.0".
- Netmask:** A text input field containing "255.255.0.0".

At the bottom right of the form, there are two buttons: "Apply" and "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 507: Showing Summary Addresses for External Routes

Routing Protocol > OSPF > Summary Address

Action: Show

Process ID: 1

Summary Address List Total: 1

	IP Address	Netmask
☐	192.168.0.0	255.255.0.0

Delete 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 711](#)).
- ◆ 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 711](#)).

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 508: Configuring Settings for All Interfaces Assigned to a VLAN

Routing Protocol > OSPF > Interface

Action: Configure by VLAN

VLAN ID: 1

Cost (1-65535): 1

Router Priority (0-255): 1

Hello Interval (1-65535): 10 sec

Dead Interval (1-65535): 40 sec

Transmit Delay (1-65535): 1 sec

Retransmit Interval (1-65535): 5 sec

Authentication Type: MD5

Message Digest Key ID: 1

Authentication Key: aiebel

Clear Click the button to clear the configuration of this VLAN.

Apply Revert

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 509: Configuring Settings for a Specific Area Assigned to a VLAN

Routing Protocol > OSPF > Interface

Action: Configure by Address

VLAN ID: 1

IP Address: 192.168.0.2

Cost (1-65535): 1

Router Priority (0-255): 1

Hello Interval (1-65535): 10 sec

Dead Interval (1-65535): 40 sec

Transmit Delay (1-65535): 1 sec

Retransmit Interval (1-65535): 5 sec

Authentication Type: MD5

Message Digest Key ID: 2

Authentication Key: bobby_martin

Clear [Click the button to clear the configuration of this IP address.](#)

Apply Revert

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 510: 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 511: Showing MD5 Authentication Keys

Routing Protocol > OSPF > Interface

Action: Show MD5 Key

VLAN ID: 1

Interface MD5 List Total: 2


	Area ID	Key ID
☐	0.0.0.0	1
☐	192.168.0.2	2

Delete 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 512: 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 732](#).

Parameters

These parameters are displayed:

- ◆ **Process ID** – Process ID as configured in the Network Area configuration screen (see [page 711](#)).
- ◆ **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 513: Adding a Virtual Link

The screenshot shows the 'Routing Protocol > OSPF > Virtual Link' configuration page. The 'Action' dropdown is set to 'Add'. Below it, the 'Process ID' is set to '1'. The 'Transit Area ID' is '192.168.10.0' and the 'Neighbor ID' is '192.168.10.3'. At the bottom right are 'Apply' and 'Revert' buttons.

To show virtual links:

1. Click Routing Protocol, OSPF, Virtual Link.
2. Select Show from the Action list.
3. Select the process ID.

Figure 514: Showing Virtual Links

The screenshot shows the 'Routing Protocol > OSPF > Virtual Link' configuration page with the 'Action' dropdown set to 'Show'. The 'Process ID' is set to '1'. Below the form is a table titled 'Virtual Link List' with 'Total: 2'.

	Transit Area ID	Neighbor ID	State	Local Address	Remote Address	Hello Due	Adjacency State
☐	0.0.0.1	10.2.0.0	Down	192.168.1.1	192.168.2.1	Inactive	Full
☐	0.0.0.2	10.3.0.0	Waiting	*	*		Down

At the bottom right are 'Delete' and 'Revert' buttons.

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 515: Configuring Detailed Settings for a Virtual Link

Routing Protocol > OSPF > Virtual Link

Action: Configure Detailed Settings

Process ID: 1

Virtual Link List Total: 1

Transit Area ID	Neighbor ID	Transmit Delay (1-65535)	Retransmit Interval (1-65535)	Hello Interval (1-65535)	Dead Interval (1-65535)	Auth Key Type	MD5 Key ID	Auth Key
192.168.10.0	192.168.10.3	1	5	10	40	Simple		Gaibel

Apply Revert

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 516: Showing MD5 Authentication Keys

Routing Protocol > OSPF > Virtual Link

Action: Show MD5 Key

Process ID: 1

Virtual Link MD5 List Total: 6

	Transit Area ID	Neighbor ID	Key ID
☐	0.0.0.1	10.2.0.0	1
☐	0.0.0.1	10.2.0.0	4
☐	0.0.0.2	10.3.0.0	3
☐	0.0.0.2	10.3.0.0	5
☐	0.0.0.2	10.3.0.0	9
☐	0.0.0.2	10.3.0.0	13

Delete Revert

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 711](#)).
- ◆ **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 517: Displaying Information in the Link State Database

Routing Protocol > OSPF > Information

Type ☒ LSDB ☐ Neighbor

Process ID

Query by :

☐ Self-Originate

☐ Link ID

☐ Adv Router

Link State Type

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 711](#)).
- ◆ **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 518: Displaying Neighbor Routers Stored in the Link State Database

The screenshot shows a web interface for OSPF configuration. At the top, the breadcrumb is "Routing Protocol > OSPF > Information". Below this, there are two radio buttons: "LSDB" (unselected) and "Neighbor" (selected). A "Process ID" dropdown menu is set to "1". Below the controls, the text "Neighbor Information List" is followed by "Total: 2". A table with five columns displays the neighbor information.

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

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 a Multicast Router”](#) on page 539).

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 519: Enabling IPv4 Multicast Routing

The screenshot shows a web interface configuration page titled "Multicast > Multicast Routing > General". Below the title bar, there is a section for "Multicast Forwarding Status" with a checkbox that is currently checked and labeled "Enabled". At the bottom right of the configuration area, there are two buttons: "Apply" and "Revert".

Web Interface (IPv6)

To enable IPv6 multicast routing:

1. Click Multicast, IPv6 Multicast Routing, General.
2. Enable Multicast Forwarding Status.
3. Click Apply.

Figure 520: 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 521: Displaying the IPv4 Multicast Routing Table

Multicast > Multicast Routing > Information					
Action: 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 522: Displaying Detailed Entries from IPv4 Multicast Routing Table

The screenshot shows the 'Multicast > Multicast Routing > Information' page. At the top, there is a breadcrumb trail and a green help icon. Below the breadcrumb, there is an 'Action:' dropdown menu set to 'Show Details'. The main content area displays the following fields:

- 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 these fields, there is a section titled 'Downstream Interface List' with a 'Total: 3' indicator. This section contains a table with two columns: 'Interface' and 'State'.

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 523: Displaying the IPv6 Multicast Routing Table

The screenshot shows the 'Multicast > IPv6 Multicast Routing > Information' page. At the top, there is a breadcrumb trail and a green help icon. Below the breadcrumb, there is an 'Action:' dropdown menu set to 'Show Summary'. The main content area displays the following fields:

- Group Address: FF02::0101
- Source Address: FE80::0101
- Interface: VLAN 4096
- Owner: PIM-DM
- Flags: Forward

Below these fields, there is a section titled 'IPv6 Multicast Routing Summary List' with a 'Total: 3' indicator. This section contains a table with five columns: 'Group Address', 'Source Address', 'Interface', 'Owner', and 'Flags'.

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 524: Displaying Detailed Entries from IPv6 Multicast Routing Table

The screenshot shows a web interface for configuring IPv6 Multicast Routing. The title bar indicates the path: Multicast > IPv6 Multicast Routing > Information. Below the title bar, there is an 'Action:' dropdown menu set to 'Show Details'. The main content area displays the following fields:

- Group Address: FF02::0101
- Source Address: FE80::0101
- Upstream Interface: VLAN 1
- Up Time: 00:00:05
- Stat Expire Time: 00:03:02
- Owner: PIM-DM
- Flags: Dense

Below these fields is a 'Query' button. At the bottom, there is a section titled 'Downstream Interface List' with a 'Total: 3' indicator. This section contains a table with two columns: 'Interface' and 'State'.

Interface	State
VLAN 1	Forward
VLAN 2	Pruned
VLAN 3	Forward

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 754](#)), 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 748](#)).

Web Interface

To enable PIM multicast routing:

1. Click Routing Protocol, PIM, General.
2. Enable PIM Routing Protocol.
3. Click Apply.

Figure 525: 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 535](#)). 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 526: Configuring PIM Interface Settings (Dense Mode)

Routing Protocol > PIM > Interface

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	☒ 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

Figure 527: 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 528: Showing PIM Neighbors

Routing Protocol > PIM > Neighbor

Neighbor Information Total: 2

Address	VLAN	Uptime	Expire	DR
192.168.0.4	1	00:01:24	00:01:21	
192.168.0.5	1	00:00:02	00:01:43	Yes

Configuring Global PIM-SM Settings

Use the Routing Protocol > 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 529: Configuring Global Settings for PIM-SM

The screenshot shows the 'Routing Protocol > PIM > SM' configuration page. The 'Step' dropdown is set to '1. Configure Global'. The configuration fields are as follows:

Field	Value	Unit/Note
Register Rate Limit (1-65535)	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)

Buttons: Apply, Revert

Configuring a PIM BSR Candidate Use the Routing Protocol > 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 763](#) and [“Configuring a PIM RP Candidate” on page 764](#)). 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 530: Configuring a PIM-SM BSR Candidate

The screenshot shows a web interface titled "Routing Protocol > PIM > SM". Below the title bar, there is a "Step:" dropdown menu set to "2. BSR Candidate". The main configuration area contains four settings: "BSR Candidate Status" with a checked checkbox and the text "Enabled"; "VLAN ID" with a dropdown menu showing "1"; "Hash Mask Length (0-32)" with a text input field containing "20"; and "Priority (0-255)" with a text input field containing "200". At the bottom right of the configuration area, there are two buttons: "Apply" and "Revert".

Configuring a PIM Static Rendezvous Point

Use the Routing Protocol > 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 531: 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 532: Showing PIM Static Rendezvous Points

Routing Protocol > PIM > SM

Step: 3. RP Address Action: Show

PIM-SM RP Address List Total: 1

	RP Address	Group Address	Group Mask
☐	192.168.1.1	224.9.0.0	255.255.255.0

Delete Revert

Configuring a PIM RP Candidate

Use the Routing Protocol > 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 533: Configuring a PIM RP Candidate

The screenshot shows the configuration page for a PIM RP candidate. The breadcrumb is 'Routing Protocol > PIM > SM'. The 'Step' dropdown is set to '4. RP Candidate' and the 'Action' dropdown is set to 'Add'. The configuration fields are as follows:

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.

Buttons: 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 534: Showing Settings for a PIM RP Candidate

The screenshot shows the configuration page for a PIM RP candidate with the 'Action' dropdown set to 'Show'. The configuration fields are as follows:

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

Buttons: Delete

Displaying the PIM BSR Router Use the Routing Protocol > 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.

Figure 535: Showing Information About the PIM BSR

Routing Protocol > PIM > SM	
Step:	5. Show Information
Action:	Show BSR Router
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

18. These parameters are based on RFC 5059.

Displaying PIM RP Mapping Use the Routing Protocol > 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 536: Showing PIM RP Mapping

Routing Protocol > PIM > SM				
Step:	5. Show Information	Action:	Show RP Mapping	
RP Mapping Information List Total: 2				
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
Clear				

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 769](#)), 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 748](#)).
- ◆ 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 537: Enabling PIMv6 Multicast Routing



The screenshot shows a web interface for configuring PIMv6. The breadcrumb navigation at the top reads "Routing Protocol > PIM6 > General". Below this, there is a section for "PIM6 Routing Protocol" with a checkbox that is checked and labeled "Enabled". At the bottom right of the configuration area, 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 device (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 effective 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 538: Configuring PIMv6 Interface Settings (Dense Mode)

Routing Protocol > PIM6 > Interface

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	☐ 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

Apply Revert

Figure 539: 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	☐ 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 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 540: Showing PIMv6 Neighbors



The screenshot shows a web interface with the breadcrumb 'Routing Protocol > PIM6 > Neighbor' and a green help icon. Below the breadcrumb is the title 'Neighbor Information' and 'Total: 2'. A table displays the following data:

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 > 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 541: Configuring Global Settings for PIM6-SM

Routing Protocol > PIM6 > SM

Step: 1. Configure Global

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)

Apply Revert

Configuring a PIM6 BSR Candidate Use the Routing Protocol > 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: 1-4094)
- ◆ **Hash Mask Length** – Hash mask length (in bits) used for RP selection (see [“Configuring a PIM6 Static Rendezvous Point” on page 778](#) and [“Configuring a PIM6 RP Candidate” on page 780](#)). 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 542: Configuring a PIM6-SM BSR Candidate

The screenshot shows a web interface titled "Routing Protocol > PIM6 > SM". Below the title bar, there is a "Step:" dropdown menu set to "2. BSR Candidate". The main configuration area contains the following fields:

- BSR Candidate Status:** A checkbox labeled "Enabled" which is checked.
- VLAN ID:** A dropdown menu showing "1".
- Hash Mask Length (0-32):** A text input field containing "10".
- Priority (0-255):** A text input field containing "255".

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

Configuring a PIM6 Static Rendezvous Point

Use the Routing Protocol > 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 543: 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, SM.
2. Select RP Address from the Step list.
3. Select Show from the Action list.

Figure 544: Showing PIM6 Static Rendezvous Points



The screenshot shows a web interface for configuring PIM6 Static Rendezvous Points. The breadcrumb trail is "Routing Protocol > PIM6 > SM". Below the breadcrumb, there are two dropdown menus: "Step: 3. RP Address" and "Action: Show". The main content area is titled "PIM6-SM RP Address List" with a "Total: 1" indicator. It contains a table with two columns: "RP Address" and "Group Address". The table has one data row with the values "FE80::200:E8FF:FE93:82A0/128" and "FF00::/8" respectively. At the bottom right of the table, there are "Delete" and "Revert" buttons.

RP Address	Group Address
FE80::200:E8FF:FE93:82A0/128	FF00::/8

Configuring a PIM6 RP Candidate

Use the Routing Protocol > 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: 1-4094)
- ◆ **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 RP is advertised for all multicast groups.
- ◆ **Group Prefix Length** – An IPv6 network prefix length for a multicast group. (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 545: Configuring a PIM6 RP Candidate

The screenshot shows a web interface titled "Routing Protocol > PIM6 > SM". Below the title bar, there are two dropdown menus: "Step:" set to "4. RP Candidate" and "Action:" set to "Add". The main configuration area contains several fields: "VLAN ID" with a dropdown menu showing "1"; "Interval (60-16383)" with a text input field and "sec (Optional)" label; "Priority (0-255)" with a text input field and "(Optional)" label; "Group Address" with a text input field containing "FFAA::0101" and "(Optional)" label; and "Group Prefix Length (8-128)" with a text input field containing "8" and "(Optional)" label. At the bottom right of the form are two buttons: "Apply" and "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 546: Showing Settings for a PIM6 RP Candidate



The screenshot shows a web interface for configuring PIM6. The breadcrumb trail is "Routing Protocol > PIM6 > SM". Below this, there are two dropdown menus: "Step: 4. RP Candidate" and "Action: Show". The main content area is titled "PIM6-SM RP Candidate List" with a "Total: 1" indicator. It contains a table with the following columns: a checkbox, "VLAN", "Interval", "Priority", and "Group Address / Prefix Length". There is one row in the table with the values: checkbox (checked), "1", "60", "0", and "FF00::/8". Below the table are "Delete" and "Revert" buttons.

	VLAN	Interval	Priority	Group Address / Prefix Length
☒	1	60	0	FF00::/8

Displaying the PIM6 BSR Router Use the Routing Protocol > 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

19. These parameters are based on RFC 5059.

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, PIM6, SM.
2. Select Show Information from the Step list.
3. Select Show BSR Router from the Action list.

Figure 547: Showing Information About the PIM6 BSR

Routing Protocol > PIM6 > SM	
Step:	5. Show Information
Action:	Show BSR Router
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 > 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 548: Showing PIM6 RP Mapping



The screenshot shows a web interface for configuring PIM6. The breadcrumb trail is "Routing Protocol > PIM6 > SM". Below this, there are two dropdown menus: "Step:" set to "5. Show Information" and "Action:" set to "Show RP Mapping". The main content area is titled "RP Mapping Information List" with a "Total: 1" indicator. It contains a table with the following data:

Groups	RP Address	Information Source	Uptime	Expire
FF00::/8	FE80::200:CFF:FE00:FD/128	static	00:24:22	Never

Below the table is a "Clear" button.

Section III

Appendices

This section provides additional information and includes these items:

- ◆ [“Software Specifications” on page 787](#)
- ◆ [“Troubleshooting” on page 793](#)
- ◆ [“License Statement / GPL Code Statement” on page 795](#)



Software Specifications

Software Features

Management Authentication Local, RADIUS, TACACS+, Port Authentication (802.1X), HTTPS, SSH, Port Security, IP Filter

General Security Measures Access Control Lists (512 rules), Port Authentication (802.1X), MAC Authentication, Port Security, DHCP Snooping, IP Source Guard

Port Configuration 1000BASE-T: 10/100 Mbps at half/full duplex, 1000 Mbps at full duplex
1000BASE-SX/LX/ZX: 1000 Mbps at full duplex (SFP+)
10GBASE-SR/LR/ER: 10 Gbps at full duplex (SFP+)

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 10 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, protocol-based, tagged (802.1Q), voice VLANs, IP subnet, MAC-based, QinQ tunnel, GVRP for automatic VLAN learning

Class of Service Supports four levels of priority
Strict, Weighted Round Robin (WRR), or a combination of strict and weighted queueing
Layer 3/4 priority mapping: IP DSCP

Quality of Service DiffServ²⁰ supports class maps, policy maps, and service policies

Multicast Filtering IGMP Snooping (Layer 2 IPv4)
MLD Snooping (Layer 2 IPv6)

IP Routing Static routes
CIDR (Classless Inter-Domain Routing)

Additional Features BOOTP Client
Connectivity Fault Management
DHCP Client, Relay, Option 82
DNS Client, Proxy
ERPS (Ethernet Ring Protection Switching)
LLDP (Link Layer Discover Protocol)
OAM (Operation, Administration, and Maintenance)
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

²⁰. Only supported for IPv4.

SNMP Management access via MIB database
Trap management to specified hosts

RMON Groups 1, 2, 3, 9 (Statistics, History, Alarm, Event)

Standards

Ethernet Service OAM (ITU-T Y.1731) - partial support
 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 802.3ac 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)
 HTTPS
 ICMP (RFC 792)
 IGMP (RFC 1112)
 IGMPv2 (RFC 2236)
 IGMPv3 (RFC 3376) - partial support
 IPv4 IGMP (RFC 3228)
 MLD Snooping (RFC 4541)
 NTP (RFC 1305)
 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)
ERPS MIB (ITU-T G.8032)
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)
NTP (RFC 1305)
P-Bridge MIB (RFC 2674P)
Port Access Entity MIB (IEEE 802.1X)
Port Access Entity Equipment MIB
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 Telnet, web browser, or SNMP software	<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.◆ If you cannot connect using Telnet, you may have exceeded the maximum number of concurrent Telnet/SSH sessions permitted. Try connecting again at a later time.
Cannot connect using Secure Shell	<ul style="list-style-type: none">◆ If you cannot connect using SSH, you may have exceeded the maximum number of concurrent Telnet/SSH sessions permitted. Try connecting again at a later time.◆ Be sure the control parameters for the SSH server are properly configured on the switch, and that the SSH client software is properly configured on the management station.◆ Be sure you have generated both an RSA and DSA public key on the switch, exported this key to the SSH client, and enabled SSH service. Try using another SSH client or check for updates to your SSH client application.◆ Be sure you have set up an account on the switch for each SSH user, including user name, authentication level, and password.◆ Be sure you have imported the client's public key to the switch (if public key authentication is used).
Cannot access the on-board configuration program via a serial port connection	<ul style="list-style-type: none">◆ Check to see if you have set the terminal emulator program to VT100 compatible, 8 data bits, 1 stop bit, no parity, and the baud rate set to 115200 bps.◆ Verify that you are using the DB-9 null-modem serial cable supplied with the switch. If you use any other cable, be sure that it conforms to the pin-out connections provided in the Installation Guide.
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 Statement / GPL Code Statement

This product resp. the here (<http://global.level1.com/downloads.php?action=init>) for downloading offered software includes software code developed by third parties, including software code subject to the GNU General Public License Version 2 ("GPLv2") and GNU Lesser General Public License 2.1 ("LGPLv2.1").

Written Offer for GPL/LGPL Source Code

We will provide everyone upon request the applicable GPLv2 and LGPLv2.1 source code files via CDROM or similar storage medium for a nominal cost to cover shipping and media charges as allowed under the GPLv2 and LGPLv2.1. This offer is valid for 3 years. GPLv2 and LGPLv2 inquiries: Please direct all GPL and LGPL inquiries to the following address:

Digital Data Communications GmbH
Zeche-Norm-Str. 25
44319 Dortmund
Deutschland

Phone: +49 231 9075 - 0
Fax: +49 231 9075 - 184

Email: support@level1.com
Web: www.level1.com

NO WARRANTY

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details. 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.

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.
- 11.** 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.

12. 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

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

one line to give the program's name and an idea of what it does.

Copyright (C) yyyy name of author

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of

MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) year name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details
type `show w'. This is free software, and you are welcome
to redistribute it under certain conditions; type `show c'
for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright
interest in the program `Gnomovision'
(which makes passes at compilers) written
by James Hacker.
```

```
signature of Ty Coon, 1 April 1989
Ty Coon, President of Vice
```

Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License.

Notification of Compliance

Europe - EU Declaration of Conformity



For complete DoC please visit

<http://global.level1.com/downloads.php?action=init>

GPL License Agreement

GPL may be included in this product, to view the GPL license agreement goes to

<http://download.level1.com/level1/gpl/GPL.pdf>

For GNU General Public License (GPL) related information, please visit

<http://global.level1.com/downloads.php?action=init>

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.

BGP Border Gateway Protocol is a protocol used to make core routing decisions on the Internet. It maintains a table of IP networks to register reachability among autonomous systems (AS). BGP makes routing decisions based on path, network policies and/or rule-sets.

BOOTP Boot Protocol. BOOTP is used to provide bootup information for network devices, including IP address information, the address of the TFTP server that contains the devices system files, and the name of the boot file.

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 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 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 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

- 802.1Q tunnel 172
 - access 180
 - configuration, guidelines 176
 - configuration, limitations 175
 - CVID to SVID map 178
 - description 172
 - ethernet type 177
 - interface configuration 179
 - mode selection 180
 - status, configuring 176
 - TPID 177
 - uplink 180
- 802.1X
 - authenticator, configuring 351
 - global settings 350
 - port authentication 348
 - port authentication accounting 281, 282
 - supplicant, configuring 356

A

- AAA
 - accounting 802.1X port settings 281, 282
 - accounting exec command privileges 281
 - accounting exec settings 282
 - accounting summary 283
 - accounting update 281
 - accounting, configuring 281
 - authorization & accounting 274
 - authorization exec settings 288
 - authorization method 288
 - authorization settings 287
 - RADIUS group settings 282
 - TACACS+ group settings 282
- acceptable frame type 166
- ACL 315
 - ARP 320, 331
 - binding to a port 332
 - IPv4 Extended 320, 323
 - IPv4 Standard 320, 322
 - IPv6 Extended 320, 327
 - IPv6 Standard 320, 325
 - MAC 320, 329
 - time range 316
- Address Resolution Protocol *See* ARP

- address table 193
 - aging time 197
 - aging time, displaying 197
 - aging time, setting 197
- ARP
 - configuration 670
 - description 670
 - proxy 670
 - statistics 674
- ARP ACL 331
- ARP inspection 336
 - ACL filter 339
 - additional validation criteria 338
 - ARP ACL 340
 - enabling globally 338
 - trusted ports 341
- ATC 230
 - control response 233
 - functional limitations 231
 - limiting traffic rates 230
 - shutting down a port 231
 - thresholds 233
 - timers 231
 - usage 230
- authentication
 - MAC address authentication 295
 - MAC, configuring ports 298
 - network access 295
 - public key 310
 - web 292
 - web authentication for ports, configuring 294
 - web authentication port information, displaying 294
 - web authentication, re-authenticating address 294
 - web authentication, re-authenticating ports 294
 - web, configuring 293
- Automatic Traffic Control *See* ATC

B

- BOOTP 614
- bootstrap router
 - PIM-SM 761
 - PIMv6-SM 777
- BPDU 204
 - filter 217
 - flooding when STA disabled on VLAN 208

Index

- flooding when STA globally disabled 208
- selecting protocol based on message format 217
- shut down port on receipt 216
- bridge extension capabilities, displaying 83
- broadcast storm, threshold 228, 229

C

- cable diagnostics 134
- canonical format indicator 248
- CFM
 - basic operations 482
 - continuity check errors 515
 - continuity check messages 471, 480, 482, 483
 - cross-check message 480, 483
 - cross-check start delay 484
 - delay measure 504
 - description 480
 - domain service access point 480, 492, 496
 - fault isolation 480
 - fault notification 480, 514
 - fault notification generator 483, 489, 514
 - fault verification 480
 - link trace cache 513
 - link trace message 480, 482, 500
 - loop back messages 480, 482, 502
 - maintenance association 480, 492
 - maintenance domain 480, 481, 487
 - maintenance end point 481, 483, 488, 492, 496, 506, 507
 - maintenance intermediate point 481, 488, 509
 - maintenance level 481, 482
 - maintenance point 480
 - MEP archive 490
 - MEP direction 497
 - remote maintenance end point 483, 492, 498, 507, 510, 511
 - service instance 480, 482
 - SNMP traps 485
- class map
 - DiffServ 252
- Class of Service *See* CoS
- clustering switches, management access 450
- committed burst size, QoS policy 260, 261
- committed information rate, QoS policy 260, 261
- community string 425
- configuration files, restoring defaults 85
- configuration settings
 - restoring 87, 88
 - saving 87
- Connectivity Fault Management *See* CFM
- continuity check errors, CFM 515
- continuity check messages, CFM 471, 480, 482, 483
- CoS 237
 - configuring 237

- default mapping to internal values 248
- enabling 244
- layer 3/4 priorities 244
- priorities, mapping to internal values 248
- queue mapping 241
- queue mode 238
- queue weights, assigning 239
- CoS/CFI to PHB/drop precedence 248
- CPU
 - status 106
 - utilization, showing 106
- cross-check message, CFM 480, 483
- cross-check start delay, CFM 484
- CVLAN to SPVLAN map 178

D

- default IPv6 gateway, configuration 619
- default priority, ingress port 237
- default settings, system 47
- delay measure, CFM 504
- designated router
 - PIM 755
 - PIMv6 771
- DHCP 614, 643
 - address pool 649
 - class identifier 644
 - client 614
 - client identifier 643, 644
 - option 82 information 377
 - relay service 645
 - relay service, enabling 646
 - server 646
- DHCPv4 snooping 375
 - circuit ID, setting 380
 - enabling 377
 - global configuration 377
 - information option 378
 - information option policy 378
 - information option, enabling 378
 - information option, remote ID 378
 - policy selection 378
 - specifying trusted interfaces 380
 - verifying MAC addresses 377
 - VLAN configuration 379
- Differentiated Code Point Service *See* DSCP
- Differentiated Services *See* DiffServ
- DiffServ 251
 - binding policy to interface 265
 - class map 252
 - classifying QoS traffic 252
 - color aware, srTCM 260
 - color aware, trTCM 261
 - color blind, srTCM 260
 - color blind, trTCM 261

- committed burst size 260, 261, 262
 - committed information rate 260, 261
 - configuring 251
 - excess burst size 261
 - metering, configuring 256, 257
 - peak burst size 262
 - peak information rate 262
 - policy map 256
 - policy map, description 253
 - QoS policy 256
 - service policy 265
 - setting CoS for matching packets 259
 - setting IP DSCP for matching packets 260, 261, 262
 - setting PHB for matching packets 259
 - single-rate, three-color meter 256, 260
 - srTCM metering 256, 260
 - traffic between CIR and BE, configuring response 260
 - traffic between CIR and PIR, configuring response 261
 - trTCM metering 261
 - two-rate, three-color meter 257
 - violating traffic, configuring response 262
- DNS**
- default domain name 637
 - displaying the cache 642
 - domain name list 637
 - enabling lookup 637
 - name server list 637
 - static entries, IPv4 641
- Domain Name Service *See* DNS
- domain service access point, CFM 480, 492, 496
- downloading software 85
- automatically 89
 - using FTP or TFTP 89
- DR priority, PIM-SM 757
- DR priority, PIMv6-SM 773
- drop precedence
- CoS priority mapping 248
 - DSCP ingress map 246
- DSA encryption 312, 313
- DSCP 244
- enabling 244
 - ingress map, drop precedence 246
 - mapping to internal values 245
- DSCP to PHB/drop precedence 246
- dynamic addresses
- clearing 199
 - displaying 197
- dynamic QoS assignment 296, 299
- dynamic VLAN assignment 295, 299
- E**
- ECMP, maximum paths 678
- edge port, STA 215, 219
- encryption
- DSA 312, 313
 - RSA 312, 313
- engine ID 414, 415
- ERPS
- configuration guidelines 458
 - control VLAN 463
 - domain configuration 460
 - domain, enabling 462
 - global configuration 459
 - guard timer 472
 - hold-off timer 471
 - major domain 468
 - MEG level 463
 - node identifier 468
 - propagate topology change 470
 - ring configuration 460
 - ring, enabling 462
 - status, displaying 475
 - wait-to-restore timer 472
 - WTR timer 472
- event logging 384
- excess burst size, QoS policy 260
- exec command privileges, accounting 281
- exec settings
- accounting 282
 - authorization 288
- F**
- fault isolation, CFM 480
- fault notification generator, CFM 483, 489, 514
- fault notification, CFM 480, 514
- fault verification, CFM 480
- firmware
- displaying version 81
 - upgrading 85
 - upgrading automatically 89
 - upgrading with FTP or TFTP 89
 - version, displaying 81
- G**
- GARP VLAN Registration Protocol *See* GVRP
- gateway, IPv6 default 619
- general security measures 273
- GVRP
- enabling 169
 - global setting 169
 - interface configuration 170
- H**
- hardware version, displaying 81
- hash mask length, PIM-SM BSR 762

Index

hash mask length, PIMv6-SM BSR 777

hello holdtime

PIM 755

PIMv6 770

hello interval

PIM 755

PIMv6 771

HTTPS 305, 306

configuring 305

replacing SSL certificate 306

secure-site certificate 306

HTTPS, secure server 305

I

IEEE 802.1D 203

IEEE 802.1s 203

IEEE 802.1w 203

IEEE 802.1X 348

IGMP

enabling per interface 573

filter profiles, configuration 557, 559

filter, parameters 557, 559

filtering & throttling 556

filtering & throttling, creating profile 557

filtering & throttling, enabling 556

filtering & throttling, interface configuration 559

filtering & throttling, status 556

groups, displaying 543, 577

interface configuration 573

interface status, displaying 575

last member query interval 574

Layer 2 533

Layer 3 569

maximum response time 574

multicast groups, displaying 577

proxy 570

proxy routing 569

proxy routing, configuring 570

proxy routing, interface configuration 573

query 535

query interval 574

report delay 574

robustness value 573

robustness variable 573

services, displaying 577

showing groups 577

snooping 533

snooping & query, parameters 535

snooping, configuring 535

snooping, immediate leave 545

static groups, configuring 575

version 573

IGMP proxy

configuration steps 571

enabling 572

unsolicited report interval 572

IGMP services, displaying 551

IGMP snooping

configuring 544

enabling per interface 544, 545

forwarding entries 551

immediate leave, status 545

interface attached to multicast router 541, 543

last leave 535

last member query interval 547

proxy address 548

proxy reporting 546

querier timeout 538

query interval 547

query response interval 547

query suppression 535

router port expire time 538

static host interface 534

static multicast routing 539

static port assignment 541

static router interface 534

static router port, configuring 539

statistics, displaying 552

TCN flood 536

unregistered data flooding 537

version exclusive 538

version for interface, setting 547

version, setting 538

with proxy reporting 535

immediate leave, IGMP snooping 545

immediate leave, MLD snooping 563

importing user public keys 313

ingress filtering 166

IP address

BOOTP/DHCP 614

setting 613

IP filter, for management access 344

IP routing 663, 691

configuring interfaces 666

maximum paths 678

unicast protocols 666

IP source guard

ACL table, learning mode 364

configuring static entries 365

learning mode, ACL table or MAC table 364

MAC table, learning mode 364

setting filter criteria 363

setting maximum bindings 364

IP statistics 629

IPv4 address

BOOTP/DHCP 614

setting 613

- IPv6
 - displaying neighbors 628
 - duplicate address detection 628
 - enabling 620
 - MTU 620
 - router advertisements, blocking 622
- IPv6 address
 - dynamic configuration (global unicast) 625
 - dynamic configuration (link-local) 620
 - EUI format 625
 - EUI-64 setting 625
 - explicit configuration 620
 - global unicast 625
 - link-local 626
 - manual configuration (global unicast) 625
 - manual configuration (link-local) 626
 - setting 618
- IPv6 source guard
 - configuring static entries 371
 - setting filter criteria 369
 - setting maximum bindings 370
- J**
- jumbo frame 82
- K**
- key
 - private 308
 - public 308
 - user public, importing 313
- key pair
 - host 308
 - host, generating 312
- L**
- LACP
 - configuration 139
 - group attributes, configuring 143
 - group members, configuring 141
 - load balancing 149
 - local parameters 146
 - partner parameters 148
 - protocol message statistics 145
 - protocol parameters 141
 - timeout, for LACPDU 140
- last member query interval, IGMP snooping 547
- license information 795
- Link Layer Discovery Protocol - Media Endpoint Discovery
 - See LLDP-MED
- Link Layer Discovery Protocol See LLDP
- link trace cache, CFM 513
- link trace message, CFM 480, 482, 500
- link type, STA 215, 219
- LLDP 389
 - device statistics details, displaying 411
 - device statistics, displaying 409
 - display device information 400
 - displaying remote information 400
 - interface attributes, configuring 391
 - message attributes 391
 - message statistics 409
 - remote information, displaying 408, 409
 - remote port information, displaying 400
 - timing attributes, configuring 389
 - TLV 389, 392
 - TLV, management address 392
 - TLV, port description 392
 - TLV, system capabilities 392
 - TLV, system description 392
 - TLV, system name 392
- LLDP-MED 389
 - notification, status 391
 - TLV 393
 - TLV, inventory 393
 - TLV, location 393
 - TLV, MED capabilities 393
 - TLV, network policy 393
- local engine ID 414
- logging
 - messages, displaying 385
 - syslog traps 386
 - to syslog servers 386
- log-in, web interface 54
- logon authentication 290
 - encryption keys 278
 - RADIUS client 277
 - RADIUS server 277
 - sequence 275
 - settings 276, 278
 - TACACS+ client 276
 - TACACS+ server 276
- loop back messages, CFM 480, 482, 502
- loopback detection
 - STA 205
- M**
- MAC address authentication 295
 - ports, configuring 298
 - reauthentication 298
- MAC address, mirroring 200
- main menu, web interface 56
- maintenance association, CFM 480, 492
- maintenance domain, CFM 480, 481, 487
- maintenance end point, CFM 481, 483, 488, 492, 496, 506, 507
- maintenance intermediate point, CFM 481, 488, 509

Index

- maintenance level, CFM 481, 482
- maintenance point, CFM 480
- management access, filtering per address 344
- management access, IP filter 344
- Management Information Bases (MIBs) 790
- matching class settings, classifying QoS traffic 253
- media-type 116
- memory
 - status 107
 - utilization, showing 107
- MEP archive, CFM 490
- mirror port
 - configuring 120
 - configuring local traffic 120
 - configuring remote traffic 122
- MLD snooping 561
 - configuring 561
 - enabling 561
 - groups, displaying 567, 568
 - immediate leave 563
 - immediate leave, status 563
 - interface attached to multicast router 564, 565
 - multicast static router port 564
 - querier 561
 - querier, enabling 561
 - query interval 562
 - query, maximum response time 562
 - robustness value 562
 - static port assignment 566
 - static router port 564
 - unknown multicast, handling 562
 - version 562
- MSTP 220
 - global settings, configuring 207, 220
 - global settings, displaying 212
 - interface settings, configuring 213, 225
 - interface settings, displaying 226
 - path cost 225
- MTU for IPv6 620
- multicast filtering 531
 - enabling IGMP snooping 545
 - enabling IGMP snooping per interface 544
 - enabling MLD snooping 561
 - router configuration 539
- multicast groups 543, 551, 567, 577
 - displaying 543, 551, 567, 577
 - static 541, 543, 566, 567
- multicast router discovery 544
- multicast router port, displaying 541, 565
- multicast routing 745
 - description 745
 - enabling, IPv4 748
 - enabling, IPv6 748
 - global settings, IPv4 748
 - global settings, IPv6 748
 - PIM 753
 - PIM-DM 753
 - PIM-SM 753, 760
 - PIMv6 769
 - PIMv6-SM 775
 - reverse path tree 746
 - routing table, IPv4 749
 - upstream interface 749
- multicast services
 - configuring 541, 566
 - displaying 543, 567
- multicast static router port 539
 - configuring 539
 - configuring for MLD snooping 564
- multicast storm, threshold 229
- Multicast VLAN Registration *See* MVR
- multicast, filtering and throttling 556
- MVR
 - assigning static multicast groups 589, 605
 - configuring 583
 - description 579
 - interface status, configuring 587
 - interface status, displaying 588
 - IP for control packets sent upstream 583
 - proxy switching 581
 - receiver groups, displaying 591
 - robust value for proxy switching 581
 - setting interface type 588
 - setting multicast domain 583, 587
 - setting multicast groups 584
 - setting multicast priority 583, 599
 - specifying a domain 583, 587
 - specifying a VLAN 583
 - specifying priority 583, 599
 - static binding 589
 - static binding, group to port 589
 - statistics, displaying 592
 - using immediate leave 588
- MVR6
 - assigning static multicast groups 605
 - configuring 599
 - interface status, configuring 603
 - interface status, displaying 605
 - IP for control packets sent upstream 599
 - proxy switching 597
 - receiver groups, displaying 607
 - robust value for proxy switching 598
 - setting interface type 604
 - setting multicast domain 604
 - setting multicast groups 600
 - specifying a domain 604
 - specifying a VLAN 599
 - static binding 605

- static binding, group to port 605
- statistics, displaying 608
- using immediate leave 605

N

network access

- authentication 295
- dynamic QoS assignment 299
- dynamic VLAN assignment 299
- MAC address filter 300
- port configuration 298
- reauthentication 298
- secure MAC information 303

NTP

- authentication keys, specifying 99
- client, enabling 95
- setting the system clock 97
- specifying servers 97

O

OAM

- active mode 517
- displaying settings and status 516
- enabling on switch ports 516
- event log, displaying 520
- message statistics, displaying 519
- mode selection 517
- passive mode 517
- remote device information, displaying 521
- remote loop back test 522
- setting to active mode 517
- setting to passive mode 517

OSPF 710

- area border router 711, 718, 720, 721, 722, 724, 725, 726, 738
- AS summary route 730
- authentication key 734
- authentication type 734
- auto cost for an interface 715
- autonomous system boundary router 716, 718, 720, 722, 725, 728, 730, 741
- backbone 711, 712, 720, 737, 738
- configuration settings, displaying 717
- cost for interface 732
- default cost for summary route 722, 724
- default external route 716
- default metric for external routes 715
- general settings 714, 717
- hello interval 733
- LSA advertisement interval 734
- LSA database, displaying 740
- message digest key 734
- neighboring router information, displaying 743

- network area 711
- normal area 712
- NSSA 719, 720, 725
- process ID 713, 714, 719, 721, 724, 725, 727, 729, 731
- process parameters, displaying 717
- redistributing external routes 728
- retransmit interval 734
- RFC 1583 compatible 714
- router ID 715
- router priority 733
- routing table, displaying 740
- SPF timers 715
- stub 719, 723
- transit area 711, 712, 721, 723, 737, 738
- transmit delay over interface 733
- virtual link 737

P

passwords 290

- administrator setting 290

path cost 219

- method 209
- STA 219

peak burst size, QoS policy 261

peak information rate, QoS policy 261

per-hop behavior, DSCP ingress map 246

PIM 753

- configuring 753
- dense-mode attributes 757
- designated router 755
- enabling for interfaces 754, 755
- enabling globally 753
- hello holdtime 755
- hello interval 755
- interface settings 754
- neighbor routers, displaying 759
- sparse-mode attributes 757

PIM-DM 753

- configuring 753
- global configuration 755–757, 760
- interface settings 757
- neighbor routers 759

PIM-SM 753, 760

- bootstrap router 761
- BSR candidate 761
- BSR elected, displaying 766
- configuring 753, 760
- DR priority 757
- global configuration 757
- hash mask length for BSR 762
- interface settings 757
- neighbor routers 759
- register rate limit for DR 760

- rendezvous point 763
 - RP candidate 764
 - RP candidate, advertising 764
 - RP mapping, displaying 768
 - shared tree 760
 - shortest path tree 760
 - SPT threshold 760
 - static RP, configuring 763
 - PIMv6 769
 - configuring 769
 - dense mode, enabling 770
 - dense-mode attributes 772
 - designated router 771
 - enabling for interfaces 769
 - enabling globally 769
 - global configuration 769
 - graft retry interval 772
 - hello holdtime 770
 - hello interval 771
 - interface configuration, displaying 774
 - interface settings 769
 - max graft retries 772
 - neighbor routers 775
 - neighbor routers, displaying 775
 - prune delay 771
 - prune state, hold time 771
 - sparse-mode attributes 773
 - state refresh message interval 772
 - triggered hello delay 772
 - PIMv6-DM
 - global configuration 770, 775
 - interface settings 772
 - PIMv6-SM 775
 - bootstrap router 777
 - BSR candidate 777
 - BSR elected, displaying 782
 - configuring 775
 - DR priority 773
 - global configuration 773
 - hash mask length for BSR 777
 - interface settings 773
 - register rate limit for DR 775
 - rendezvous point 778
 - RP candidate 780
 - RP candidate, advertising 780
 - RP mapping, displaying 783
 - shared tree 776
 - shortest path tree 776
 - SPT threshold 776
 - static RP, configuring 778
 - policing traffic, QoS policy 256, 260
 - policy map
 - DiffServ 256
 - port authentication 348
 - port priority
 - configuring 237
 - default ingress 237
 - STA 214
 - port security, configuring 346
 - port, statistics 126
 - ports
 - autonegotiation 117
 - broadcast storm threshold 228, 229
 - capabilities 117
 - configuring 116
 - duplex mode 117
 - flow control 117
 - forced selection of media type 116
 - mirroring 120
 - mirroring local traffic 120
 - mirroring remote traffic 122
 - multicast storm threshold 229
 - speed 117
 - statistics 126
 - transceiver threshold, auto-set 132
 - transceiver threshold, trap 132
 - unknown unicast storm threshold 229
 - power savings
 - configuring 151
 - enabling per port 151
 - PPPoE 657–661
 - priority, default port ingress 237
 - private key 308
 - problems, troubleshooting 793
 - protocol migration 217
 - protocol VLANs 181
 - configuring 181
 - interface configuration 183
 - system configuration 181
 - proxy address, IGMP snooping 548
 - proxy ARP 670
 - proxy reporting, IGMP snooping 546
 - public key 308
 - PVID, port native VLAN 166
- ## Q
- QinQ Tunneling *See* 802.1Q tunnel
 - QoS 251
 - configuring 251
 - CoS/CFI to PHB/drop precedence 248
 - DSCP to PHB/drop precedence 245
 - dynamic assignment 299
 - matching class settings 253
 - PHB to queue 241
 - selecting DSCP, CoS 244
 - QoS policy
 - committed burst size 260, 261
 - excess burst size 260

- peak burst size 261
- policing flow 256, 260
- srTCM 256
- srTCM police meter 260
- trTCM 257
- trTCM police meter 261
- QoS policy, committed information rate 260, 261
- QoS policy, peak information rate 261
- Quality of Service *See* QoS
- query interval, IGMP snooping 547
- query response interval, IGMP snooping 547
- queue weight, assigning to CoS 239

R

- RADIUS
 - logon authentication 277
 - settings 277
- rate limit
 - port 227
 - setting 227
- register rate limit, PIM-SM 760
- register rate limit, PIMv6-SM 775
- remote engine ID 414
- remote logging 386
- remote maintenance end point, CFM 483, 492, 498, 507, 510, 511
- rendezvous point
 - PIM-SM 763
 - PIMv6-SM 778
- restarting the system 111
 - at scheduled times 111
 - showing restart time 114
- RIP 692
 - authentication key 706
 - authentication mode 706
 - configuring 692
 - default external route 694
 - default metric 694
 - description 691
 - global settings 693
 - interface protocol settings 704
 - interface, enabling 697
 - neighbor router 700, 709
 - passively monitoring updates 699
 - poison reverse 692, 707
 - protocol packets, receiving 706
 - protocol packets, sending 706
 - receive version 706
 - redistributing external routing information 701
 - routes, displaying 709
 - routing table, clearing 696
 - send version 706
 - specifying interfaces 697
 - split horizon 692, 707

- timers 695
- version 693
- RMON 440
 - alarm, displaying settings 442
 - alarm, setting thresholds 440
 - event settings, displaying 444
 - response to alarm setting 443
 - statistics history, collection 445
 - statistics history, displaying 446
 - statistics, collection 448
 - statistics, displaying 449
- router redundancy
 - protocols 681
 - VRRP 681
- routing table, displaying 677
- RSA encryption 312, 313
- RSTP 203
 - global settings, configuring 207
 - global settings, displaying 212
 - interface settings, configuring 213
 - interface settings, displaying 218

S

- secure shell 308
 - configuration 308
- security, general measures 273
- serial port, configuring 103
- service instance, CFM 480, 482
- shared tree
 - PIM-SM 760
 - PIMv6-SM 776
- shortest path tree
 - PIM-SM 760
 - PIMv6-SM 776
- Simple Mail Transfer Protocol *See* SMTP
- Simple Network Management Protocol *See* SNMP
- single rate three color meter *See* srTCM
- SMTP
 - event handling 387
 - sending log events 387
- SNMP 411
 - community string 425
 - enabling traps 432
 - enabling traps, mac-address changes 201
 - filtering IP addresses 344
 - global settings, configuring 413
 - trap manager 432
 - traps, CFM 485
 - users, configuring 426, 429
- SNMPv3 414–433
 - engine ID 414, 415
 - engine identifier, local 414
 - engine identifier, remote 414, 415
 - groups 419

Index

- local users, configuring 426
 - remote users, configuring 429
 - user configuration 426, 429
 - views 417
 - SNTP
 - setting the system clock 95
 - specifying servers 96
 - software
 - displaying version 81
 - downloading 85
 - version, displaying 81
 - Spanning Tree Protocol *See* STA
 - specifications, software 787
 - SPT threshold, PIM-SM 760
 - SPT threshold, PIMv6-SM 776
 - srTCM
 - police meter 260
 - QoS policy 256
 - SSH 308
 - authentication retries 311
 - configuring 308
 - downloading public keys for clients 313
 - generating host key pair 312
 - server, configuring 310
 - timeout 311
 - SSL, replacing certificate 306
 - STA 203
 - BPDU filter 217
 - BPDU flooding 208, 214
 - BPDU shutdown 216
 - detecting loopbacks 205
 - edge port 215, 219
 - global settings, configuring 207
 - global settings, displaying 212
 - interface settings, configuring 213
 - interface settings, displaying 218
 - link type 215, 219
 - loopback detection 205
 - MSTP interface settings, configuring 225
 - MSTP path cost 225
 - path cost 219
 - path cost method 209
 - port priority 214
 - port/trunk loopback detection 205
 - protocol migration 217
 - transmission limit 209
 - stack
 - configuring 108
 - master unit, setting 108
 - renumbering 110
 - stacking ports, enabling 109
 - standards, IEEE 789
 - startup files
 - creating 85
 - displaying 85
 - setting 85
 - static addresses, setting 195
 - static routes, configuring 675
 - statistics
 - ARP 674
 - statistics, port 126
 - STP 207
 - summary, accounting 283
 - summer time, setting 101
 - switch clustering, for management 450
 - switch settings
 - restoring 87
 - saving 87
 - system clock
 - setting 93
 - setting manually 94
 - setting the time zone 100
 - setting with NTP 97
 - setting with SNTP 95
 - summer time 101
 - system software, downloading from server 85
- ## T
- TACACS+
 - logon authentication 276
 - settings 278
 - TCN
 - flood 536
 - general query solicitation 537
 - Telnet
 - configuring 105
 - server, enabling 105
 - time range, ACL 316
 - time zone, setting 100
 - time, setting 93
 - TPID 177
 - traffic segmentation 153
 - assigning ports 153
 - enabling 153
 - sessions, assigning ports 155
 - sessions, creating 154
 - transceiver data, displaying 131
 - transceiver thresholds
 - configuring 132
 - displaying 132
 - trap manager 432
 - troubleshooting 793, 795
 - trTCM
 - police meter 261
 - QoS policy 257
 - trunk
 - configuration 135
 - LACP 139

- static 137
- tunneling unknown VLANs, VLAN trunking 156
- two rate three color meter *See* trTCM
- Type Length Value
 - See* LLDP TLV

U

- UDLD
 - configuration 525
 - interface settings 527
 - neighbor information 529
 - protocol intervals 526
- UDP helper 653
 - application port 654
 - application server 655
 - description 653
 - destination port 654
 - enabling 654
 - forward destination 655
 - target subnet 655
 - UDP ports 654
- unicast routing 691
 - ECMP 678
 - ECMP maximum paths 678
- unidirectional link detection 525
- unknown unicast storm, threshold 229
- unregistered data flooding, IGMP snooping 537
- upgrading software 85
- upstream interface, multicast route 749
- user account 290
- user password 290

V

- VLAN trunking 156
- VLANs 159–189
 - 802.1Q tunnel mode 180
 - acceptable frame type 166
 - adding static members 165
 - creating 162
 - description 159
 - displaying port members by interface 168
 - displaying port members by interface range 169
 - displaying port members by VLAN index 168
 - dynamic assignment 299
 - egress mode 165
 - ingress filtering 166
 - interface configuration 165
 - IP subnet-based 185
 - MAC-based 187
 - mirroring 189
 - protocol 181
 - protocol, configuring 181
 - protocol, configuring groups 181

- protocol, interface configuration 183
- protocol, system configuration 181
- PVID 166
- tag swapping 191
- tunneling unknown groups 156
- voice 267
- voice VLANs 267
 - detecting VoIP devices 268
 - enabling for ports 271
 - identifying client devices 269
- VoIP traffic 267
 - ports, configuring 270
 - telephony OUI, configuring 269
 - voice VLAN, configuring 268
- VoIP, detecting devices 271
- VRRP 681
 - authentication 685
 - configuration settings 681, 682
 - group statistics 689
 - preemption 683, 685
 - priority 683, 685
 - protocol message statistics 688
 - timers 684
 - virtual address 683

W

- web authentication 292
 - address, re-authenticating 294
 - configuring 293
 - port information, displaying 294
 - ports, configuring 294
 - ports, re-authenticating 294
- web interface
 - access requirements 53
 - configuration buttons 55
 - home page 54
 - menu list 56
 - panel display 55

